Veritas™ Cluster Server Administrator's Guide

Windows 2000, Windows Server 2003

5.0 Release Update 1

symantec™

Veritas Cluster Server Administrator's Guide

Veritas Cluster Server

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

Veritas Cluster Server is a licensed product. See the *Storage Foundation and High Availability Solutions for Windows, Installation and Upgrade Guide* for license installation instructions.

Technical support

For technical assistance, visit http://entsupport.symantec.com 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 I Clustering concepts and terminology

Chapter 1 Introducing Veritas Cluster Server

What is a VCS cluster? 16
Can my application be clustered? 18
Physical components of VCS 20
Logical components of VCS 21
Putting the pieces together 31

Chapter 2 About cluster topologies

Basic failover configurations 34
Advanced failover configurations 38
Cluster topologies and storage configurations 41

Chapter 3 VCS configuration concepts

About configuring VCS 48
About the VCS configuration language 48
About the main.cf file 49
The types.cf file 51
About VCS attributes 52
About VCS keywords and reserved words 57
VCS environment variables 58

Section II Administration-Putting VCS to work

Chapter 4 About the VCS user privilege model

About VCS user privileges and roles 64
How administrators assign roles to users 68
User privileges for OS user groups in secure clusters 68
About VCS privileges for users with multiple roles 70

Chapter 5 Getting started with VCS

Prerequisites 72

Configuring cluster components using the VCS configuration wizard 73
Configuring a cluster from the command line 89
Analyzing a domain 95
Administering a cluster 100
Administering service groups in a cluster 102
Administering resources in a cluster 118
Administering resource types in a cluster 128
Administering systems in a cluster 129
Administering attributes in a cluster 131
Conducting a search 133

Chapter 6 Administering the cluster from Cluster Manager (Java console)

About the Cluster Manager (Java Console) 136
Getting started 136
Reviewing components of the Java Console 138
About Cluster Monitor 140
About Cluster Explorer 146
Accessing additional features of the Java Console 159
Administering Cluster Monitor 168
Administering user profiles 173
Administering service groups 178
Administering resources 200
Administering systems 219
Administering clusters 222
Executing commands 223
Editing attributes 223
Querying the cluster configuration 226
Setting up VCS event notification using the Notifier wizard 227
Administering logs 230
Administering VCS Simulator 234

Chapter 7 Administering the cluster from the command line

About administering VCS from the command line 236
Starting VCS 237
Stopping VCS 238
Managing VCS configuration files 241
Managing VCS users from the command line 249
Querying VCS 252
Administering service groups 259
Administering agents 264

Administering resources 264
Administering systems 272
Administering clusters 273
Administering resource types 274
Using the -wait option in scripts 275
Administering simulated clusters from the command line 276

Chapter 8 Configuring resources and applications in VCS

Configuring shared storage 278
Configuring network resources 283
Configuring file shares 288
Configuring print shares 300
Configuring Microsoft virtual machines 313
Configuring IIS sites 317
Configuring services 326
Configuring processes 330
Configuring the MSMQ resource 332
Configuring infrastructure and support agents 335
Configuring applications using the Application Configuration Wizard 343

Chapter 9 Modifying the cluster configuration

Adding nodes to a cluster 372
Removing nodes from a cluster 380
Reconfiguring a cluster 384
Enabling and disabling Symantec Product Authentication Service 390
Configuring the ClusterService group 392
Deleting a cluster configuration 401

Chapter 10 Predicting VCS behavior using VCS Simulator

About VCS Simulator 406
Installing VCS Simulator 406
Administering VCS Simulator from the Java Console 408
Administering VCS Simulator from the command line 413

Section III VCS communication and operations

Chapter 11 About communications, membership, and data protection in the cluster

About cluster communications 420
About cluster membership 424

About membership arbitration 427
About data protection 431
Examples of VCS operation with I/O fencing 432
About cluster membership and data protection without I/O fencing 437
Examples of VCS operation without I/O fencing 439
Summary of best practices for cluster communications 444

Chapter 12 Controlling VCS behavior

About VCS behavior on resource faults 448
Controlling VCS behavior at the service group level 451
Controlling VCS behavior at the resource level 459
Changing agent file paths and binaries 471
Service group workload management 472
Sample configurations depicting workload management 475

Chapter 13 The role of service group dependencies

About service group dependencies 492
Service group dependency configurations 497
Group Dependency FAQs 505
Linking service groups 506
VCS behavior with service group dependencies 506

Section IV Administration-Beyond the basics

Chapter 14 VCS event notification

About VCS event notification 512
Components of VCS event notification 515
VCS events and traps 517
Monitoring aggregate events 525
Detecting complementary events 526
Configuring notification 526

Chapter 15 VCS event triggers

About VCS event triggers 528
Using event triggers 528
List of event triggers 529

Section V Multi-cluster configurations

Chapter 16 Connecting clusters–Creating global clusters

How VCS global clusters work 544
VCS global clusters: The building blocks 545
Prerequisites for global clusters 551
Setting up a global cluster 553
When a cluster faults 563
Setting up a fire drill 565

Chapter 17 Administering DR configurations

About creating global service groups 568
Administering global clusters 568
Administering global service groups 574
Administering global heartbeats 578
Running a fire drill 580
Viewing fire drill logs 581

Chapter 18 Administering global clusters from Cluster Manager (Java console)

About global clusters 584
Adding a remote cluster 585
Deleting a remote cluster 589
Administering global service groups 592
Administering global heartbeats 596

Chapter 19 Administering global clusters from the command line

About administering global clusters from the command line 602
Global querying 602
Administering global service groups 609
Administering resources 611
Administering clusters in global clusters 612
Administering heartbeats 614

Chapter 20 Setting up replicated data clusters

About replicated data clusters 616
How VCS replicated data clusters work 617
Setting up a replicated data cluster configuration 618

Section VI Troubleshooting and performance

Chapter 21 VCS performance considerations

How cluster components affect performance 622

How cluster operations affect performance 626
Monitoring CPU usage 632
VCS agent statistics 633
VCS performance with non-HA products 635
VCS performance with SFW 636

Chapter 22 Troubleshooting and recovery for VCS

Logging 637
Handling network failure 640
Troubleshooting VCS startup 647
Troubleshooting secure clusters 656
Troubleshooting service groups 656
Troubleshooting resources 659
Troubleshooting notification 660
Troubleshooting Cluster Management Console (Single Cluster Mode) or Web Console 661
Troubleshooting VCS configuration backup and restore 666
Troubleshooting and recovery for global clusters 667
VCS utilities 671

Section VII Appendixes

Appendix A VCS user privileges—administration matrices

About administration matrices 685
Administration matrices 686

Appendix B Cluster and system states

Remote cluster states 694
System states 696

Appendix C VCS attributes

About attributes and their definitions 702
Resource attributes 702
Resource type attributes 708
Service group attributes 715
System attributes 728
Cluster attributes 735
Heartbeat attributes 742

Appendix D Administering Symantec Web Server

About Symantec Web Server 746
Getting Started 746
Configuring ports for VRTSweb 748
Managing VRTSweb SSL certificates 752
Configuring SMTP notification for VRTSweb 755
Configuring logging for VRTSweb 761
Modifying the maximum heap size for VRTSweb 766

Appendix E Accessibility and VCS

About accessibility in VCS 768
Navigation and keyboard shortcuts 768
Support for accessibility settings 769
Support for assistive technologies 769

Appendix F Configuring LLT over UDP

When to use LLT over UDP 771
Configuring LLT over UDP 771

Appendix G Handling concurrency violation in Any-to-Any configurations

Concurrency violation scenario 779
About the vcsgensvc.vbs script 780
Sample configuration to handle concurrency violation 781

Glossary 783

Index 787

Section I

Clustering concepts and terminology

- Chapter 1, “Introducing Veritas Cluster Server” on page 15
- Chapter 2, “About cluster topologies” on page 33
- Chapter 3, “VCS configuration concepts” on page 47

Chapter 1

Introducing Veritas Cluster Server

- What is a VCS cluster?
- Can my application be clustered?
- Physical components of VCS
- Logical components of VCS
- Putting the pieces together

What is a VCS cluster?

Veritas Cluster Server (VCS) from Symantec connects multiple, independent systems into a management framework for increased availability. Each system, or node, runs its own operating system and cooperates at the software level to form a cluster. VCS links commodity hardware with intelligent software to provide application failover and control. When a node or a monitored application fails, other nodes can take predefined actions to take over and bring up services elsewhere in the cluster.

How VCS detects failure

VCS detects failure of an application by issuing specific commands, tests, or scripts to monitor the overall health of an application. VCS also determines the health of underlying resources supporting the application, such as file systems and network interfaces.

VCS uses a redundant network heartbeat to discriminate between the loss of a system and the loss of communication between systems. VCS also uses SCSI3-based membership coordination and data protection for detecting failure on a node and on fencing.

See "About cluster control, communications, and membership" on page 27.

How VCS ensures application availability

When VCS detects an application or node failure, VCS brings application services up on a different node in a cluster. VCS virtualizes IP addresses and system names, so client systems continue to access the application and are unaware of which server they use.

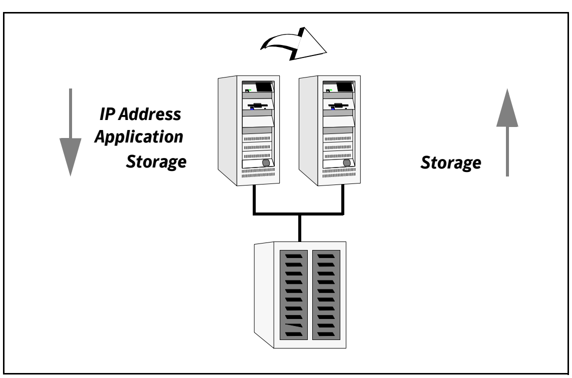

For example, in a 2-node cluster consisting of db-server1 and db-server2, a virtual address may be called db-server. Clients access db-server and are unaware of which physical server hosts the db-server.

About switchover and failover

Switchover and failover are the processes of bringing up application services on a different node in a cluster.

Switchover	A switchover is an orderly shutdown of an application and its supporting resources on one server and a controlled startup on another server.
Failover	A failover is similar to a switchover, except the ordered shutdown of applications on the original node may not be possible, so the services are started on another node.

Can my application be clustered?

Most applications can be placed under cluster control provided the following guidelines are met:

- Defined start, stop, and monitor procedures.
- Ability to restart in a known state.
- Ability to store required data on shared disks.
- Adherence to license requirements and host name dependencies.

Defined start, stop, and monitor procedures

The application to be clustered must have defined procedures for starting, stopping, and monitoring.

Start procedure	The application must have a command to start it and all resources it may require. VCS brings up the required resources in a specific order, then brings up the application using the defined start procedure. For example, to start an Oracle database, VCS must know which Oracle utility to call, such as sqlplus. VCS must also know the Oracle user, instance ID, Oracle home directory, and the pfile.
Stop procedure	An individual instance of the application must be capable of being stopped without affecting other instances. For example, killing all HTTPd processes on a Web server is unacceptable because it would also stop other Web servers. If VCS cannot stop an application cleanly, it may call for a more forceful method, like a kill signal. After a forced stop, a clean-up procedure may be required for various process- and application-specific items that may be left behind, such as shared memory segments or semaphores.
Monitor procedure	The application must have a monitor procedure that determines if the specified application instance is healthy. The application must allow individual monitoring of unique instances. For example, the monitor procedure for a Web server connects to the specified server and verifies that it is serves Web pages. In a database environment, the monitoring application can connect to the database server and perform SQL commands to verify read and write access to the database. The closer a test comes to matching what a user does, the better the test is in discovering problems. You should balance the level of monitoring between ensuring the application is up and minimizing monitor overhead.

Ability to restart the application in a known state

When the application is taken offline, it must close out all tasks, store data properly on shared disk, and exit. Stateful servers must not keep that state of clients in memory. States should be written to shared storage to ensure proper failover.

Commercial databases such as Oracle, Sybase, or SQL Server are good examples of well-written, crash-tolerant applications. On any client SQL request, the client is responsible for holding the request until it receives acknowledgement from the server. When the server receives a request, it is placed in a special *redo* log file. The data is confirmed as being saved before acknowledging the client. After a server crashes, the database recovers to the last-known committed state by mounting the data tables and applying the redo logs. This returns the database to the time of the crash. The client resubmits any outstanding client requests that are unacknowledged by the server, and all others are contained in the redo logs.

If an application cannot recover gracefully after a server crashes, it cannot run in a cluster environment. The takeover server cannot start up because of data corruption and other problems.

External data storage

The application must be capable of storing all required data and configuration information on shared disks. The exception to this rule is a true *shared nothing* cluster.

See “Shared nothing cluster” on page 43.

For example, SQLServer 2000 can be set up so that the binaries are installed on the local system, but the shared database and configuration information are stored on a shared disk.

The application must also store data to disk rather than maintaining it in memory. The takeover system must be capable of accessing all required information. This requirement precludes the use of anything inside a single system inaccessible by the peer. NVRAM accelerator boards and other disk-caching mechanisms for performance are acceptable, but must be done on the external array and not on the local host.

Licensing and host name issues

The application must be capable of running on all servers that are designated as potential hosts. This means strict adherence to licensing requirements and host name dependencies. Changing host names can lead to significant management issues when multiple systems have the same host name after an outage. Custom scripting to modify a system host name on failover is not recommended.

Symantec recommends you configure applications and licensing to run properly on all hosts.

Physical components of VCS

A VCS cluster comprises of systems that are connected with a dedicated communications infrastructure. VCS refers to a system that is part of a cluster as a node.

Each cluster has a unique cluster ID. Systems in a cluster are connected by redundant cluster communication links.

Nodes

VCS nodes host the service groups (managed applications). Each system is connected to networking hardware, and usually also to storage hardware. The systems contain components to provide resilient management of the applications, and start and stop agents.

Nodes can be individual systems, or they can be created with domains or partitions on enterprise-class systems. Individual cluster nodes each run their own operating system and possess their own boot device. Each node must run the same operating system within a single VCS cluster.

Clusters can have from 1 to 32 nodes. Applications can be configured to run on specific nodes within the cluster.

Shared storage

Storage is a key resource of most applications services, and therefore most service groups. A managed application can only be started on a system that has access to its associated data files. Therefore, a service group can only run on all systems in the cluster if the storage is shared across all systems. In many configurations, a storage area network (SAN) provides this requirement.

I/O fencing technology can be used for data protection to block access to shared storage from any system that is not a current and verified member of the cluster.

See "About cluster topologies" on page 33.

Networking

Networking in the cluster is used for the following purposes:

- Communications between the cluster nodes and the customer systems.
- Communications between the cluster nodes.

Logical components of VCS

VCS is comprised of several components that provide the infrastructure to cluster an application.

Resources and resource dependencies

Resources are hardware or software entities that make up the application. Resources include disk groups and file systems, network interface cards (NIC), IP addresses, and applications.

Resource dependencies indicate resources that depend on each other because of application or operating system requirements. Resource dependencies are graphically depicted in a hierarchy, also called a tree, where the resources higher up (parent) depend on the resources lower down (child).

Figure 1-1 shows the hierarchy for a database application.

Figure 1-1 Sample resource dependency graph

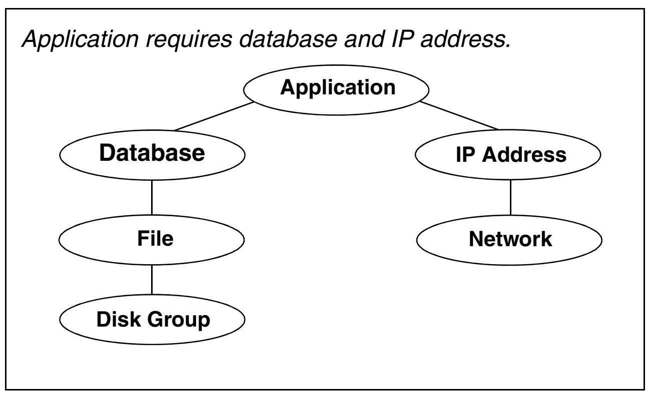

Resource dependencies determine the order in which resources are brought online or taken offline. For example, a disk group must be imported before volumes in the disk group start, and volumes must start before file systems are mounted. Conversely, file systems must be unmounted before volumes stop, and volumes must stop before disk groups are deported.

A parent is brought online after each child is brought online, and so on up the tree, until finally the application is started. Conversely, to take a managed application offline, you stop resources beginning at the top of the hierarchy. In this example, the application is stopped first, followed by the database application. Next the IP address and file systems can be stopped concurrently, since they do not have any resource dependency between them, and so on down the tree.

Child resources must be online before parent resources can be brought online. Parent resources must be taken offline before child resources can be taken offline. If resources do not have parent-child interdependencies, they can be brought online or taken offline concurrently.

Categories of resources

Different types of resources require different levels of control. In VCS there are three categories of resources:

- **On-Off.** VCS starts and stops On-Off resources as required. For example, VCS imports a disk group when required, and deports it when it is no longer needed.
- **On-Only.** VCS starts On-Only resources, but does not stop them.
  For example, in the case of the FileOnOnly resource, VCS creates the file, but does not delete the file if the service group is taken offline.
- **Persistent.** These resources cannot be brought online or taken offline. For example, a network interface card cannot be started or stopped, but it is required to configure an IP address. A Persistent resource has an operation value of None. VCS monitors Persistent resources to ensure their status and operation. Failure of a Persistent resource triggers a service group failover.

Resource types

VCS defines a resource type for each resource it manages. For example, the NIC resource type can be configured to manage network interface cards. Similarly, all IP addresses can be configured using the IP resource type.

VCS includes a set of predefined resources types. For each resource type, VCS has a corresponding agent, which provides the logic to control resources.

See "About agents in VCS" on page 25.

Service groups

A service group is a virtual container that contains all the hardware and software resources that are required to run the managed application. Service groups allow VCS to control all the hardware and software resources of the managed application as a single unit. When a failover occurs, resources do not fail over individually— the entire service group fails over. If there is more than one service group on a system, a group may fail over without affecting the others.

Figure 1-2 Typical database service group

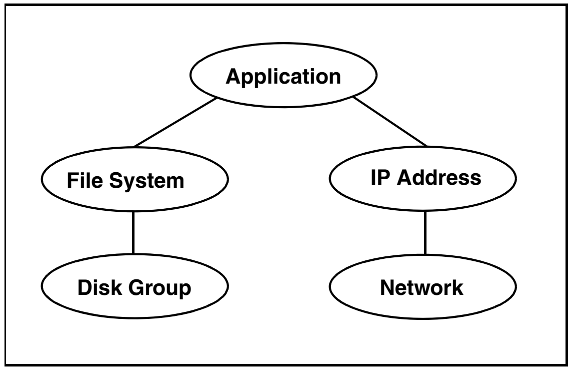

A single node may host any number of service groups, each providing a discrete service to networked clients. If the server crashes, all service groups on that node must be failed over elsewhere.

Service groups can be dependent on each other. For example a finance application may be dependent on a database application. Because the managed application consists of all components that are required to provide the service, service group dependencies create more complex managed applications. When using service group dependencies, the managed application is the entire dependency tree.

See "About service group dependencies" on page 492.

Types of service groups

VCS service groups fall in three main categories: *failover*, *parallel*, and *hybrid*.

Failover service groups

A failover service group runs on one system in the cluster at a time. Failover groups are used for most applications that do not support multiple systems to simultaneously access the application's data.

Parallel service groups

A parallel service group runs simultaneously on more than one system in the cluster. A parallel service group is more complex than a failover group. Parallel service groups are appropriate for applications that manage multiple application instances running simultaneously without data corruption.

Hybrid service groups

A hybrid service group is for replicated data clusters and is a combination of the failover and parallel service groups. It behaves as a failover group *within* a system zone and a parallel group *across* system zones.

See "About system zones" on page 452.

A hybrid service group cannot fail over across system zones. VCS allows a switch operation on a hybrid group only if both systems are within the same system zone. If there are no systems within a zone to which a hybrid group can fail over, VCS invokes the nofailover trigger on the lowest numbered node. Hybrid service groups adhere to the same rules governing group dependencies as do parallel groups.

See "About service group dependencies" on page 492.

See "nofailover event trigger" on page 531.

About the ClusterService group

The ClusterService group is a special purpose service group, which contains resources that are required by VCS components. The group contains resources for:

- Cluster Management Console
- Notification
- wide-area connector (WAC) process, which is used in global clusters

The ClusterService group can fail over to any node despite restrictions such as frozen. The ClusterService group the first service group to come online and cannot be autodisabled. The group comes online on the first node that goes in the running state. The VCS engine discourages taking the group offline manually.

About agents in VCS

Agents are multi-threaded processes that provide the logic to manage resources. VCS has one agent per resource type. The agent monitors all resources of that type; for example, a single IP agent manages all IP resources.

When the agent is started, it obtains the necessary configuration information from VCS. It then periodically monitors the resources, and updates VCS with the resource status.

The action to bring a resource online or take it offline differs significantly for each resource type. For example, bringing a disk group online requires importing the disk group. But, bringing a database online requires starting the database manager process and issuing the appropriate startup commands.

VCS monitors resources when they are online *and* offline to ensure they are not started on systems on which they are not supposed to run. For this reason, VCS starts the agent for any resource that is configured to run on a system when the cluster is started. If no resources of a particular type are configured, the agent is not started. For example, if there are no Oracle resources in your configuration, the Oracle agent is not started on the system.

Agent operations

Agents carry out specific operations on resources. The functions an agent performs are called *entry points*. For details on any of the following entry points, see the *Veritas Cluster Server Agent Developer's Guide*.

- Online—Brings a specific resource ONLINE from an OFFLINE state.
- Offline—Takes a resource from an ONLINE state to an OFFLINE state.
- Monitor—Tests the status of a resource to determine if the resource is online or offline. The entry point runs at the following times:
  - During initial node startup, to probe and determine status of all resources on the system.
  - After every online and offline operation.
  - Periodically, to verify that the resource remains in its correct state. Under normal circumstances, the monitor is run every 60 seconds when a resource is online, and every 300 seconds when a resource is expected to be offline.
- Clean—Cleans up after a resource fails to come online, fails to go offline, or fails while in an ONLINE state. The clean entry point is designed to clean up after an application. The entry point ensures that the host system is returned to a valid state. For example, the clean function may remove shared memory segments or IPC resources that are left behind by a database.

- Action—Performs actions that can be completed in a short time and which are outside the scope of traditional activities such as online and offline. Some agents have predefined action scripts that you can run by invoking the action entry point.
- Info—Retrieves specific information for an online resource.
  The retrieved information is stored in the resource attribute ResourceInfo. This entry point is invoked periodically by the agent framework when the resource type attribute InfoInterval is set to a non-zero value. The InfoInterval attribute indicates the period after which the info entry point must be invoked. For example, the Mount agent may use this entry point to indicate the space available on the file system.

Agent classifications

Bundled agents

Bundled agents are packaged with VCS. They include agents for Disk, Mount, IP, and various other resource types. See the *Veritas Cluster Server Bundled Agents Reference Guide* for a complete list.

Enterprise agents

Enterprise agents control third-party applications and are licensed separately. These include agents for Oracle, Sybase, and DB2. Contact your sales representative for more information.

Custom agents

Custom agents can be developed by you or by Symantec consultants. Typically, agents are developed because the user requires control of an application that the current bundled or enterprise agents do not support. See the *Veritas Cluster Server Agent Developer's Guide* for information on developing a custom agent.

About the VCS agent framework

The VCS agent framework is a set of common, predefined functions that are compiled into each agent. These functions include the ability to connect to the VCS engine (HAD) and to understand common configuration attributes. The agent framework frees the developer from developing support functions that are required by the cluster, and instead focus on controlling a specific resource type. For more information on developing agents, see the *Veritas Cluster Server Agent Developer's Guide*.

About cluster control, communications, and membership

Cluster communications ensure that VCS is continuously aware of the status of each system's service groups and resources. They also enable VCS to recognize which systems are active members of the cluster, which have joined or left the cluster, and which have failed.

About the high-availability daemon (HAD)

The VCS high-availability daemon (HAD) runs on each system. Also known as the VCS engine, HAD is responsible for:

- building the running cluster configuration from the configuration files
- distributing the information when new nodes join the cluster
- responding to operator input
- taking corrective action when something fails.

The engine uses agents to monitor and manage resources. It collects information about resource states from the agents on the local system and forwards it to all cluster members.

The local engine also receives information from the other cluster members to update its view of the cluster. HAD operates as a *replicated state machine* (RSM). running on each node has a completely synchronized view of the resource status on each node. Each instance of HAD follows the same code path for corrective action, as required.

The RSM is maintained through the use of a purpose-built communications package consisting of the protocols *Low Latency Transport* (LLT) and *Group Membership Services/Atomic Broadcas*t (GAB).

The hashadow process monitors HAD and restarts it when required.

About Group Membership Services/Atomic Broadcast (GAB)

The Group Membership Services/Atomic Broadcast protocol (GAB) is responsible for cluster membership and cluster communications.

- Cluster Membership
  GAB maintains cluster membership by receiving input on the status of the heartbeat from each node via LLT. When a system no longer receives heartbeats from a peer, it marks the peer as DOWN and excludes the peer from the cluster. In VCS, memberships are sets of systems participating in the cluster. There are different types of memberships.
  - A regular membership includes systems that communicate with each other across one or more network channels.

- A jeopardy membership includes systems that have only one private communication links.
- Daemon Down Node Alive (DDNA) is a condition in which the VCS high availability daemon (HAD) on a node fails, but the node is running. In a DDNA condition, VCS does not have information about the state of service groups on the node. So, VCS places all service groups that were online on the affected node in the autodisabled state. The service groups that were online on the node cannot fail over. Manual intervention is required to enable failover of autodisabled service groups. The administrator must release the resources running on the affected node, clear resource faults, and bring the service groups online on another node.

- Cluster Communications
  GAB's second function is reliable cluster communications. GAB provides guaranteed delivery of point-to-point and broadcast messages to all nodes. The VCS engine uses a private IOCTL (provided by GAB) to tell GAB that it is alive.

About Low Latency Transport (LLT)

VCS uses private network communications between cluster nodes for cluster maintenance. The Low Latency Transport functions as a high-performance, low-latency replacement for the IP stack, and is used for all cluster communications. Symantec recommends two independent networks between all cluster nodes, which provide the required redundancy in the communication path and enable VCS to discriminate between a network failure and a system failure. LLT has two major functions.

- Traffic Distribution
  LLT distributes (load balances) internode communication across all available private network links. This distribution means that all cluster communications are evenly distributed across all private network links (maximum eight) for performance and fault resilience. If a link fails, traffic is redirected to the remaining links.
- Heartbeat
  LLT is responsible for sending and receiving heartbeat traffic over network links. This heartbeat is used by the Group Membership Services function of GAB to determine cluster membership.

About security services

VCS uses the Symantec Product Authentication Service to provide secure communication between cluster nodes and clients, including the Java and the Web consoles. VCS uses digital certificates for authentication and uses SSL to encrypt communication over the public network.

In secure mode:

- VCS uses platform-based authentication.
- VCS does not store user passwords.
- All VCS users are system and domain users and are configured using fully-qualified user names. For example, administrator@vcsdomain. VCS provides a single sign-on mechanism, so authenticated users need not sign on each time to connect to a cluster.

VCS requires a system in your enterprise to be configured as a *root broker*. Additionally, all nodes in the cluster must be configured as *authentication brokers*.

- A root broker serves as the main registration and certification authority; it has a self-signed certificate and can authenticate other brokers. The root broker may be a system in the cluster, but the recommended practice is to have a single root broker per domain, typically a data center, acting as root broker for all products using Symantec Product Authentication Services. The root broker is only used during initial creation of an authentication broker.
- Authentication brokers serve as intermediate registration and certification authorities. An authentication broker authenticates users, such as a login to the cluster, or services, such as daemons running on application nodes. but cannot authenticate other brokers. Authentication brokers have certificates that are signed by the root. Each node in VCS serves as an authentication broker.
  Security credentials for the authentication broker must be obtained from the root broker.

For secure communication, VCS components acquire credentials from the authentication broker that is configured on the local system. The acquired certificate is used during authentication and is presented to clients for the SSL handshake. VCS and its components specify the account name and the domain in the following format:

- **HAD Account**
  ```
  name = _HA_VCS_(systemname)
  domain = HA_SERVICES@(fully_qualified_system_name)
  ```
- **CmdServer**

```
name = _CMDSERVER_VCS_(systemname)
domain = HA_SERVICES@(fully_qualified_system_name)
```

- Wide-area connector

```
name = _WAC_GCO_(systemname)
domain = HA_SERVICES@(fully_qualified_system_name)
```

Components for administering VCS

VCS provides the following components to administer clusters:

Cluster Management Console

A Web-based graphical user interface for monitoring and administering the cluster.

- Install the Cluster Management Console on cluster nodes to manage a single cluster.
  See "Administering the cluster from the Cluster Management Console" on page 99.
- Install the Cluster Management Console on a management server outside the cluster to manage multiple clusters.
  See the *Veritas Cluster Management Console Implementation Guide* fore more information.

Cluster Manager (Java console)

A cross-platform Java-based graphical user interface that provides complete administration capabilities for your cluster. The console runs on any system inside or outside the cluster, on any operating system that supports Java.

See"Administering the cluster from Cluster Manager (Java console)" on page 135.

VCS command-line interface (CLI)

The VCS command-line interface provides a comprehensive set of commands for managing and administering the cluster.

See "Administering the cluster from the command line" on page 235.

Putting the pieces together

In the following example, a two-node cluster shares directories to clients. Both nodes are connected to shared storage, which enables them access to the directories being shared. A single service group, "FileShare_Group," is configured to fail over between System A and System B. The service group consists of various resources, each with a different resource type.

The VCS engine, HAD, reads the configuration file, determines what agents are required to control the resources in the service group, and starts the agents. HAD determines the order in which to bring the resources online, based on the resource dependencies. VCS issues online commands to the corresponding agents in the correct order.

The following figure shows the dependency graph for the service group FileShare_Group.

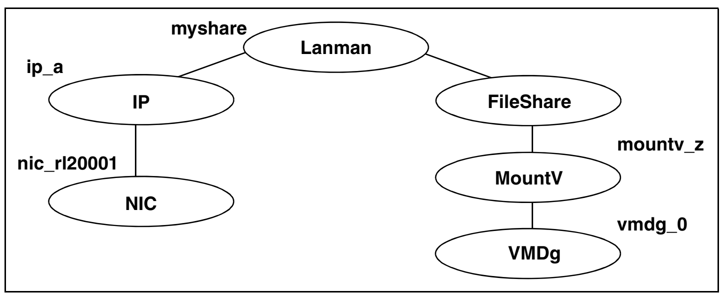

In this configuration, HAD will start agents for the disk group, mount, share, NIC, and IP on all systems configured to run FileShare_Group. The resource dependencies are configured as:

- The MountV resource requires the VMDg resource to be online before it can be brought online. The FileShare resource requires the MountV resource to be online before it can be brought online.
- The IP resource requires the NIC resource to be online before it can be brought online. The NIC resource is a persistent resource and does not need to be started.
- The Lanman resource required the FileShare and IP resources to be online before it can be brought online.

The service group can be configured to start automatically on either node in the preceding example. It can then move or fail over to the second node on command or automatically if the first node fails. Upon failover or relocation, VCS will offline the resources beginning at the top of the graph and start them on the second node beginning at the bottom.

About cluster topologies

- Basic failover configurations
- Advanced failover configurations
- Cluster topologies and storage configurations

Basic failover configurations

This section describes basic failover configurations, including asymmetric, symmetric, and N-to-1.

Asymmetric or Active/Passive configuration

In an asymmetric configuration, an application runs on a primary, or master, server. A dedicated redundant server is present to take over on any failure. The redundant server is not configured to perform any other functions. In the following illustration, a database application is moved, or failed over, from the master to the redundant server.

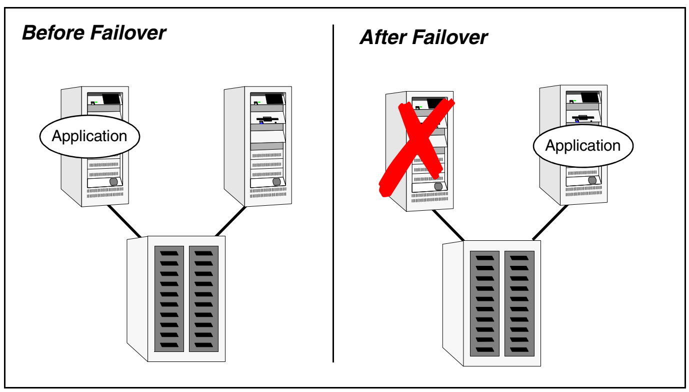

Figure 2-1 Asymmetric failover

This configuration is the simplest and most reliable. The redundant server is on stand-by with full performance capability. If other applications are running, they present no compatibility issues.

Symmetric or Active/Active configuration

In a symmetric configuration, each server is configured to run a specific application or service and provide redundancy for its peer. In this example, each server runs one application service group. When a failure occurs, the surviving server hosts both application groups.

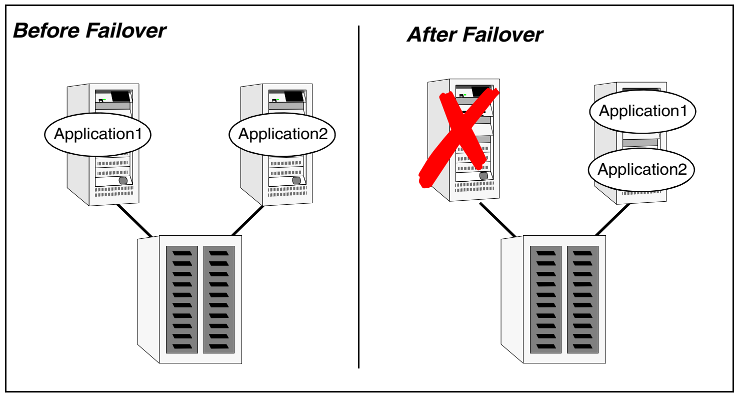

Figure 2-2 Symmetric failover

Symmetric configurations appear more efficient in terms of hardware utilization. In the asymmetric example, the redundant server requires only as much processor power as its peer. On failover, performance remains the same. In the symmetric example, the redundant server requires adequate processor power to run the existing application and the new application it takes over.

Further issues can arise in symmetric configurations when multiple applications running on the same system do not co-exist properly. Some applications work well with multiple copies started on the same system, but others fail. Issues can also arise when two applications with different I/O and memory requirements on the same system.

N-to-1 configuration

An N-to-1 failover configuration reduces the cost of hardware redundancy and still provides a potential, dedicated spare. In an asymmetric configuration there is no performance penalty and there no issues with multiple applications running on the same system; however, the drawback is the 100 percent redundancy cost at the server level.

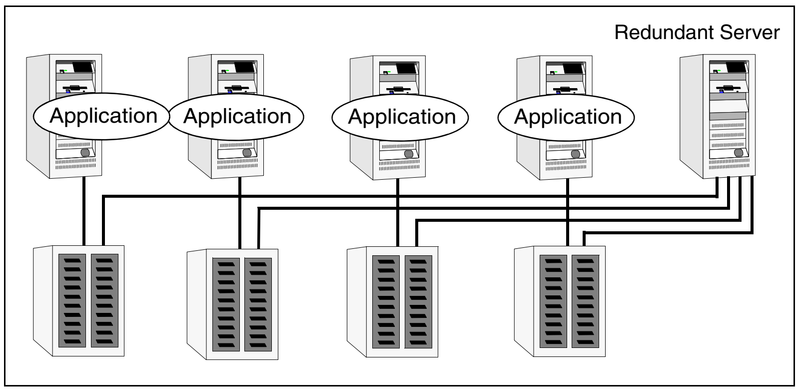

Figure 2-3 N-to-1 configuration

An N-to-1 configuration is based on the concept that multiple, simultaneous server failures are unlikely; therefore, a single redundant server can protect multiple active servers. When a server fails, its applications move to the redundant server. For example, in a 4-to-1 configuration, one server can protect four servers, which reduces redundancy cost at the server level from 100 percent to 25 percent. In this configuration, a dedicated, redundant server is cabled to all storage and acts as a spare when a failure occurs.

The problem with this design is the issue of *failback*. When the failed server is repaired, all services hosted on the server must be failed back to the server. The failback action frees the spare server and restores redundancy to the cluster.

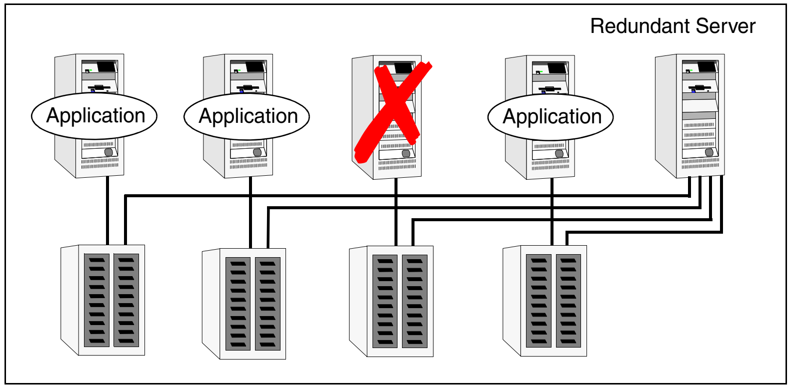

Figure 2-4 N-to-1 failover requiring failback

Most shortcomings of early N-to-1 cluster configurations were the limitations of storage architecture. Typically, it was impossible to connect more than two hosts to a storage array without complex cabling schemes and their inherent reliability problems, or resorting to expensive arrays with multiple controller ports.

Advanced failover configurations

The advent of SANs and second-generation high-availability (HA) products such as VCS, has enabled several new and useful failover configurations, described in the following sections.

N + 1 configuration

With the capabilities introduced by storage area networks (SANs), you can not only create larger clusters, but more importantly, can connect multiple servers to the same storage.

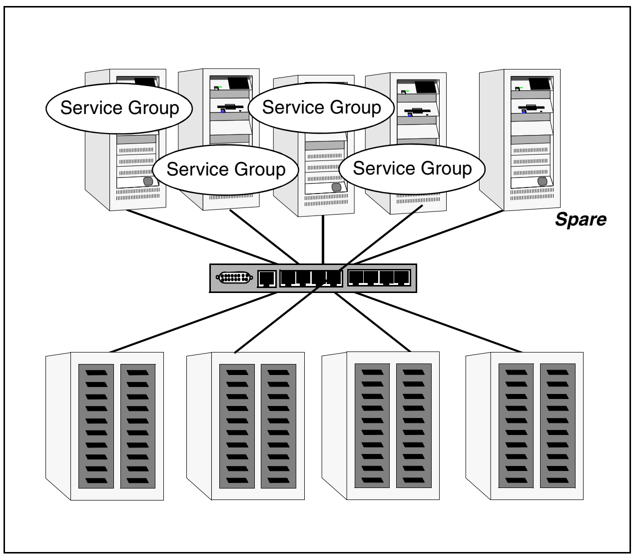

Figure 2-5 N+1 configuration

A dedicated, redundant server is no longer required in the configuration. Instead of N-to-1 configurations, you can use an *N+1* configuration. In advanced N+1 configurations, an extra server in the cluster is spare capacity only.

When a server fails, the application service group restarts on the spare. After the server is repaired, it becomes the spare. This configuration eliminates the need for a second application failure to fail back the service group to the primary system. Any server can provide redundancy to any other server.

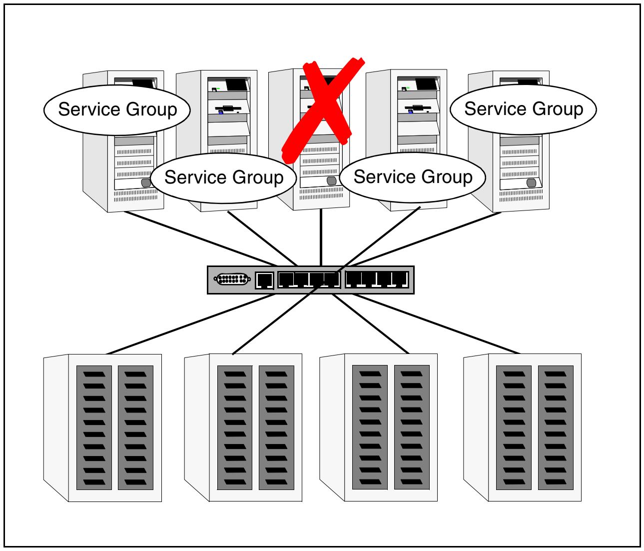

Figure 2-6 N+1 failover

N-to-N configuration

An N-to-N configuration refers to multiple service groups running on multiple servers, with each service group capable of being failed over to different servers. For example, consider a four-node cluster with each node supporting three critical database instances.

Figure 2-7 N-to-N configuration

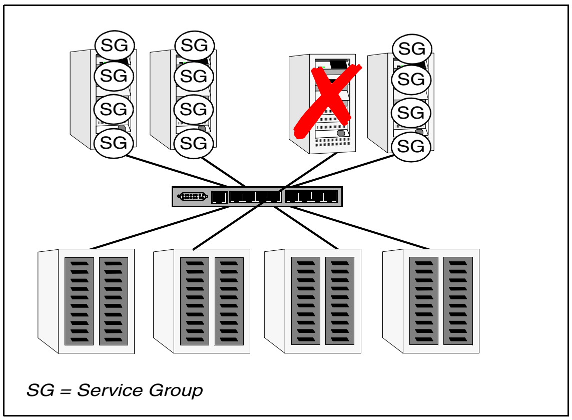

If any node fails, each instance is started on a different node, ensuring no single node becomes overloaded. This configuration is a logical evolution of N + 1: it provides cluster *standby capacity* instead of a *standby server*.

N-to-N configurations require careful testing to ensure all applications are compatible. Applications must also have complete control of where service groups fail when an event occurs.

Cluster topologies and storage configurations

This section describes commonly-used cluster topologies, along with the storage configuration used to support the topologies.

Basic shared storage cluster

In this configuration, a single cluster shares access to a storage device, typically over a SAN. An application can only be started on a node with access to the required storage. For example, in a multi-node cluster, any node designated to run a specific database instance must have access to the storage where the database's tablespaces, redo logs and control files are stored. Shared disk architecture is also the easiest to implement and maintain. When a node or application fails, all data to start on another node is stored on the shared disk.

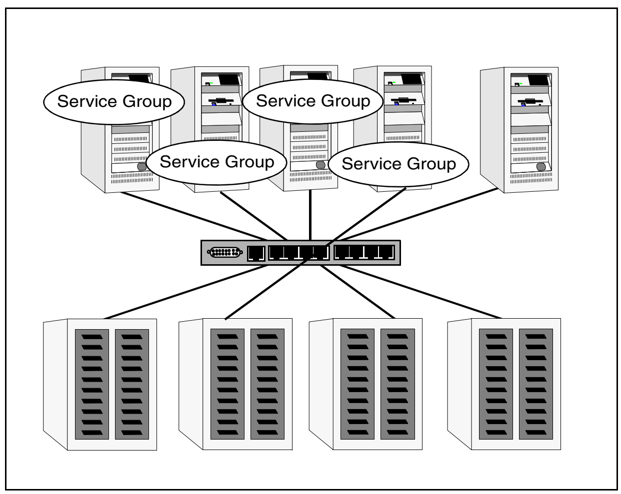

Figure 2-8 Shared disk architecture for basic cluster

Campus, or Metropolitan, shared storage cluster

In a campus environment, VCS and Veritas Volume Manager are used to create a cluster that spans multiple data centers or buildings. Instead of a single storage array, data is mirrored between arrays using Veritas Volume Manager. This provides synchronized copies of data at both sites. This procedure is identical to mirroring between two arrays in a data center, only now it is spread over a distance.

Figure 2-9 Campus shared storage cluster

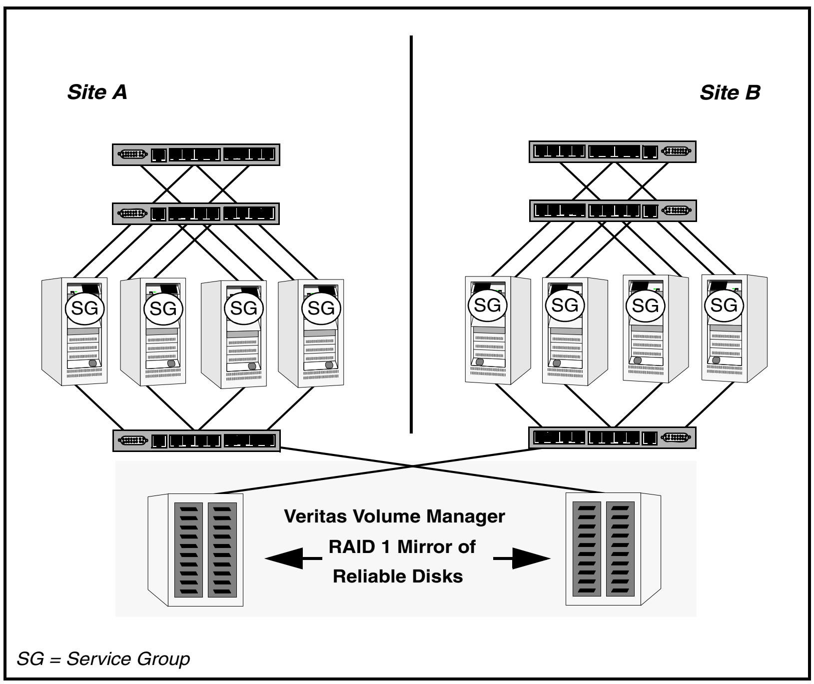

A campus cluster requires two independent network links for heartbeat, public network connectivity between buildings on same IP subnet, and two storage arrays, each providing highly available disks.

Shared nothing cluster

Systems in shared nothing clusters do not share access to disks; they maintain separate copies of data. VCS shared nothing clusters typically have read-only data stored locally on both systems. For example, a pair of systems in a cluster that includes a critical Web server, which provides access to a backend database. The Web server runs on local disks and does not require data sharing at the Web server level.

Figure 2-10 Shared nothing cluster

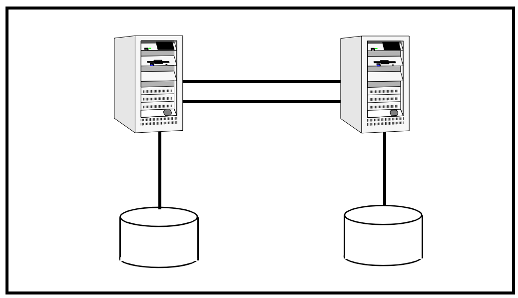

Replicated data cluster

In a replicated data cluster there no shared disk. Instead, a data replication product synchronizes copies of data between nodes. Replication can take place at the application, host, and storage levels. Application-level replication products, such as Oracle DataGuard, maintain consistent copies of data between systems at the SQL or database levels. Host-based replication products, such as Veritas Volume Replicator, maintain consistent storage at the logical volume level. Storage- or array-based replication maintains consistent copies of data at the disk or RAID LUN level.

The following illustration shows a hybrid shared storage/replicated data cluster, in which different failover priorities are assigned to nodes according to particular service groups.

Figure 2-11 Shared storage replicated data cluster

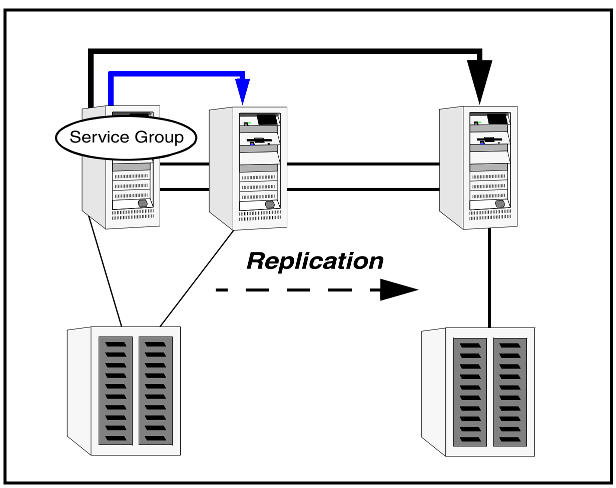

Replicated data clusters can also be configured without the ability to fail over locally, but this configuration is not recommended.

Global cluster

A global cluster links clusters at separate locations and enables wide-area failover and disaster recovery.

Local clustering provides local failover for each site or building. Campus and replicated cluster configurations offer protection against disasters affecting limited geographic regions. Large scale disasters such as major floods, hurricanes, and earthquakes can cause outages for an entire city or region. In such situations, you can ensure data availability by migrating applications to sites located considerable distances apart.

Figure 2-12 Global cluster

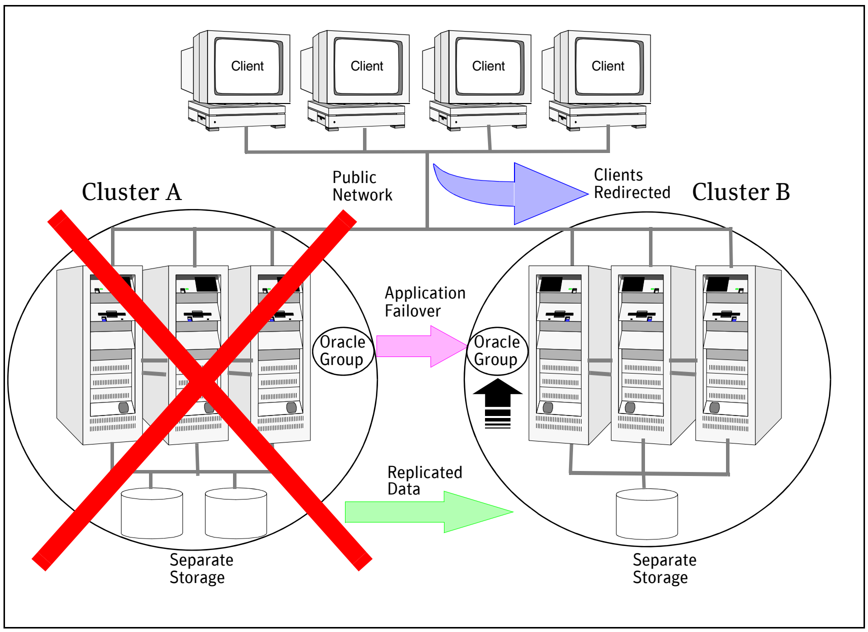

In a global cluster, if an application or a system fails, the application is migrated to another system within the same cluster. If the entire cluster fails, the application is migrated to a system in another cluster. Clustering on a global level also requires replicating shared data to the remote site.

See "How VCS global clusters work" on page 544.

Chapter 3

VCS configuration concepts

- About the VCS configuration language
- About the main.cf file
- The types.cf file
- About VCS attributes
- About VCS keywords and reserved words
- VCS environment variables

About configuring VCS

Configuring VCS means conveying to the VCS engine the definitions of the cluster, service groups, resources, and resource dependencies. VCS uses two configuration files in a default configuration:

- main.cf—Defines the cluster, including services groups and resources.
- types.cf—Defines the resource types.

By default, both files reside in the following directory:

`%VCS_HOME%\conf\config`

Additional files similar to types.cf may be present if agents have been added, such as Oracletypes.cf.

In a VCS cluster, the first system to be brought online reads the configuration file and creates an internal (in-memory) representation of the configuration. Systems brought online after the first system derive their information from systems running in the cluster.

You must stop the cluster while you are modifying the files from the command line. Changes made by editing the configuration files take effect when the cluster is restarted. The node on which the changes were made should be the first node to be brought back online.

About the VCS configuration language

The VCS configuration language specifies the makeup of service groups and their associated entities, such as resource types, resources, and attributes. These specifications are expressed in configuration files, whose names contain the suffix .cf.
There are several ways to generate configuration files:

- Use the web-based Cluster Management Console
- Use Cluster Manager (Java Console).
- Use the command-line interface.
- If VCS is not running, use a text editor to create and modify the files.

About the main.cf file

The format of the main.cf file comprises include clauses and definitions for the cluster, systems, service groups, and resources. The main.cf file also includes service group and resource dependency clauses.

- Include clauses—Include clauses incorporate additional configuration files into main.cf. These additional files typically contain type definitions, including the types.cf file. Typically, custom agents add type definitions in their own files.

  ```
  include "types.cf"
  ```

- Cluster definition—Defines the attributes of the cluster, including the cluster name and the names of the cluster users.

  ```
  cluster demo (
  UserNames = { admin = cDRpdxPmHzpS }
  )"
  ```

  See “Cluster attributes” on page 735.

- System definition—Lists the systems designated as part of the cluster. The system names must match the name returned by the command `uname -a`. Each service group can be configured to run on a subset of systems defined in this section.

  ```
  system Server1
  system Server2
  ```

- Service group definition—Service group definitions in main.cf comprise the attributes of a particular service group.

  ```
  group FileShare_Group (
      SystemList = { SystemA, SystemB }
      AutoStartList = { SystemA }
  )
  ```

  See “Service group attributes” on page 715.
  See “About the SystemList attribute.” on page 50.

- Resource definition—Defines each resource used in a particular service group. Resources can be added in any order and the utility hacf arranges the resources alphabetically the first time the configuration file is run.

  ```
  NIC NIC_resource (
      MACAddress @ system1= "02-B0-D0-D1-88-0E"
      MACAddress @ system2= "50-B0-D0-D1-88-23"
  )
  ```

- Resource dependency clause—Defines a relationship between resources. A dependency is indicated by the keyword requires between two resource names.

  ```
  IP_resource requires NIC_resource
  ```

  See “Resources and resource dependencies” on page 21.

- Service group dependency clause—To configure a service group dependency, place the keyword requires in the service group declaration of the main.cf file. Position the dependency clause before the resource dependency specifications and after the resource declarations.

  ```
  group_x requires group_y
  ```

  See "About service group dependencies" on page 492.

About the SystemList attribute.

The SystemList attribute designates all systems on which a service group can come online. By default, the order of systems in the list defines the priority of systems used in a failover. For example, the following definition configures SystemA to be the first choice on failover, followed by SystemB and then SystemC.

```
SystemList = { SystemA, SystemB, SystemC}
```

System priority may also be assigned explicitly in the SystemList attribute by assigning numeric values to each system name. For example:

```
SystemList = {SystemA=0, SystemB=1, SystemC=2}
```

If you do not assign numeric priority values, VCS assigns a priority to the system without a number by adding 1 to the priority of the preceding system. For example, if the SystemList is defined as follows, VCS assigns the values SystemA = 0, SystemB = 2, SystemC = 3.

```
SystemList = {SystemA, SystemB=2, SystemC},
```

Note that a duplicate numeric priority value may be assigned in some situations:

```
SystemList = {SystemA, SystemB=0, SystemC}
```

The numeric values assigned are SystemA = 0, SystemB = 0, SystemC = 1.

To avoid this situation, do not assign any numbers or assign different numbers to each system in SystemList.

Initial configuration

When VCS is installed, a basic main.cf configuration file is created with the cluster name, systems in the cluster, and a Cluster Manager user named *admin* with the password *password*.

The following is an example of the main.cf for cluster demo and systems SystemA and SystemB.

```
include "types.cf"
cluster demo (
UserNames = { admin = cDRpdxPmHzpS }
)
system SystemA
system SystemB
```

The types.cf file

The types.cf file describes standard resource types to the VCS engine; specifically, the data required to control a specific resource.

```
type IP (
    static i18nstr ArgList[] = { Address, SubNetMask,
        MACAddress}
    str Address
    str SubNetMask
    str MACAddress
)
```

The types definition performs two important functions:

- Defines the type of values that may be set for each attribute.
  In the IP example, the Address attribute is classified as str, or string.
  See"Attribute data types" on page 53.
- Defines the parameters passed to the VCS engine through the ArgList attribute. The line static str ArgList[] = { xxx, yyy, zzz } defines the order in which parameters are passed to the agents for starting, stopping, and monitoring resources.

For another example, review the following main.cf and types.cf representing an IP resource:

main.cf for Windows

```
IP IP_resource (
    Address = "192.168.1.201"
    SubNetMask = "255.255.254.0"
    MACAddress @ system1= "02-B0-D5-D1-88-0E"
    MACAddress @ system2= "04-B0-D0-D1-88-43"
)
```

types.cf for Windows

```
type IP (
    static i18nstr ArgList[] = { Address, SubNetMask,
    MACAddress}
    str Address
    str SubNetMask
    str MACAddress
)
```

The high-availability address is configured on the interface defined by the Address attribute.

The IP address is enclosed in double quotes because the string contains periods.

See "Attribute data types" on page 53.

The VCS engine passes the identical arguments to the IP agent for online, offline, clean and monitor. It is up to the agent to use the arguments it requires. All resource names must be unique in a VCS cluster.

About VCS attributes

VCS components are configured using *attributes*. Attributes contain data about the cluster, systems, service groups, resources, resource types, agent, and heartbeats if using global clusters. For example, the value of a service group's SystemList attribute specifies on which systems the group is configured and the priority of each system within the group. Each attribute has a definition and a value. Attributes also have default values assigned when a value is not specified.

Attribute data types

VCS supports the following data types for attributes.

String	A string is a sequence of characters enclosed by double quotes. A string may also contain double quotes, but the quotes must be immediately preceded by a backslash. A backslash is represented in a string as \\. Quotes are not required if a string begins with a letter, and contains only letters, numbers, dashes (-), and underscores (_). For example, a string defining a network interface such as hme0 or eth0 does not require quotes as it contains only letters and numbers. However a string defining an IP address contains periods and requires quotes, such as: 192.168.100.1 VCS also supports UTF-8 encoded values for some attributes. See “Localizable attributes” on page 56.
Integer	Signed integer constants are a sequence of digits from 0 to 9. They may be preceded by a dash, and are interpreted in base 10. Integers cannot exceed the value of a 32-bit signed integer: 21471183247.
Boolean	A boolean is an integer, the possible values of which are 0 (false) and 1 (true).

Attribute dimensions

VCS attributes have the following dimensions.

Scalar	A scalar has only one value. This is the default dimension.
Vector	A vector is an ordered list of values. Each value is indexed using a positive integer beginning with zero. Use a comma (,) or a semi-colon (;) to separate values. A set of brackets ([]) after the attribute name denotes that the dimension is a vector. For example, an agent's ArgList is defined as: `static str ArgList[] = { Address, SubNetMask, MACAddress }`
Keylist	A keylist is an unordered list of strings, and each string is unique within the list. Use a comma (,) or a semi-colon (;) to separate values. For example, to designate the list of systems on which a service group will be started with VCS (usually at system boot): `AutoStartList = {SystemA; SystemB; SystemC}`
Association	An association is an unordered list of name-value pairs. Use a comma (,) or a semi-colon (;) to separate values. A set of braces ({}) after the attribute name denotes that an attribute is an association. For example, to designate the list of systems on which the service group is configured to run and the system's priorities: `SystemList = {SystemA=1, SystemB=2, SystemC=3}`

Attributes and cluster objects

VCS has the following types of attributes, depending on the cluster object the attribute applies to.

Cluster attributes	Attributes that define the cluster. For example, ClusterName and ClusterAddress.
Service group attributes	Attributes that define a service group in the cluster. For example, Administrators and ClusterList.
System attributes	Attributes that define the system in the cluster. For example, Capacity and Limits.
Resource type attributes	Attributes that define the resource types in VCS. These can be further classified as: ■ **Type-independent**—Attributes that all agents (or resource types) understand. Examples: RestartLimit and MonitorInterval; these can be set for any resource type. Typically, these attributes are set for all resources of a specific type. For example, setting MonitorInterval for the IP resource type affects all IP resources. ■ **Type-dependent**—Attributes that apply to a particular resource type. These attributes appear in the type definition file (types.cf) for the agent. Examples: The MountPath attribute applies only to the Mount resource type. The Address attribute applies only to the IP resource type. Attributes defined in the file types.cf apply to all resources of a particular resource type. Defining these attributes in the main.cf file overrides the values in the types.cf file for a specific resource. For example, setting StartVolumes = 1 for the DiskGroup types.cf defaults StartVolumes to True for all DiskGroup resources. Setting the value in main.cf overrides the value on a per-resource basis. ■ **Static**—These attributes apply for every resource of a particular type. These attributes are prefixed with the term static and are not included in the resource's argument list. You can override some static attributes and assign them resource-specific values. See "Overriding resource type static attributes" on page 274.
Resource attributes	Attributes that define a specific resource. Some of these attributes are type-independent. For example, you can configure the Critical attribute for any resource. Some resource attributes are type-dependent. For example, the Address attribute defines the IP address associated with the IP resource. These attributes are defined in the main.cf file.

Localizable attributes

VCS supports UTF-8 encoded localized values for some attributes. These attributes are identified by the i18nstr keyword in the type definition file types.cf.

For example, in the FileOnOff agent, the attribute PathName is a localizable attribute.

```
type FileOnOff (
    static i18nstr ArgList[] = { PathName }
    i18nstr PathName
)
```

You can add a localizable string attribute using the `haattr -add -i18nstring` command.

Attribute scope across systems: global and local attributes

An attribute whose value applies to all systems is *global* in scope. An attribute whose value applies on a per-system basis is *local* in scope. The at operator (@) indicates the system to which a local value applies.

An example of local attributes can be found in the IP resource type where Mac addresses and routing options are assigned per machine.

```
IP IP_resource (
Address = "192.168.1.201"
SubNetMask = "255.255.254.0"
MACAddress @ system1= "02-B1-D5-D1-88-0E"
MACAddress @ system2= "04-B0-D0-D1-88-43"
)
```

Attribute life: temporary attributes

You can define temporary attributes in the types.cf file. The values of temporary attributes remain in memory as long as the VCS engine (HAD) is running. Values of temporary attributes are not available when HAD is restarted. These attribute values are not stored in the main.cf file.

Temporary attributes cannot be converted to permanent, and vice-versa. When you save a configuration, VCS saves temporary attributes and their default values in the file types.cf.

The scope of these attributes can be local to a node or global across all nodes in the cluster. Local attributes can be defined even when the node is not part of a cluster.

You can define and modify these attributes only while VCS is running.

See "Adding, deleting, and modifying resource attributes" on page 266.

About VCS keywords and reserved words

The following list includes the current keywords reserved for the VCS configuration language. Note they are case-sensitive.

action	false	local	requires	stop
after	firm	offline	resource	str
ArgListValues	global	online	set	system
before	group	MonitorOnly	Signaled	System
boolean	Group	Name	soft	temp
cluster	hard	NameRule	start	type
Cluster	heartbeat	Path	Start	Type
condition	int	Probed	state	
ConfidenceLevel	IState	remote	State	
event	keylist	remotecluster	static	

VCS environment variables

Table 3-1 lists VCS environment variables.

Table 3-1 VCS environment variables

Environment Variable	Definition and Default Value
PERL5LIB	Root directory for Perl executables. **Default:** Install Drive:\Program Files\VERITAS\cluster server\lib\perl5.
VCS_CONF	Root directory for VCS configuration files. **Default:** Install Drive:\Program Files\VERITAS\cluster server\conf\config **Note:** If this variable is added or modified you must reboot the system to apply the changes.
VCS_DOMAIN	The Security domain in which users are configured.
VCS_DOMAINTYPE	Type of domain: unixpwd, nt, nis, nisplus, or vx.
VCS_DIAG	Directory where VCS dumps HAD cores.
VCS_ENABLE_LDF	Designates whether or not log data files (LDFs) are generated. If set to 1, LDFs are generated. If set to 0, they are not.
VCS_HOME	Root directory for VCS executables. **Default:** Install Drive:\Program Files\VERITAS\cluster server\
VCS_HOST	VCS node on which `ha` commands will be run.
VCS_GAB_PORT	GAB port to which VCS connects. **Default:** h
VCS_GAB_TIMEOUT	Timeout in milliseconds for HAD to send heartbeats to GAB. **Default:** 15000 **Note:** If the specified timeout is exceeded, GAB kills HAD, and all active service groups on system are disabled.
VCS_HAD_RESTART_TIMEOUT	Set this variable to designate the amount of time the hashadow process waits (sleep time) before restarting HAD. **Default:** 0

Table 3-1 VCS environment variables

Environment Variable	Definition and Default Value
VCS_LOG	Root directory for log files and temporary files. **Default:** Install Drive:\Program Files\VERITAS\cluster server\ **Note:** If this variable is added or modified you must reboot the system to apply the changes.
VCS_SERVICE	Name of configured VCS service. **Default:** vcs **Note:** The specified service should be configured before starting the VCS engine (HAD). If a service is not specified, the VCS engine starts with port 14141.
VCS_TEMP_DIR	Directory in which temporary information required by, or generated by, hacf is stored. **Default:** Install Drive:\Program Files\VERITAS\cluster server\ This directory is created in /tmp under the following conditions: ■ The variable is not set. ■ The variable is set but the directory to which it is set does not exist. ■ The utility hacf cannot find the default location.

Section

II

Administration-Putting VCS to work

- Chapter 4, "About the VCS user privilege model" on page 63
- Chapter 5, "Getting started with VCS" on page 71
- Chapter 7, "Administering the cluster from Cluster Manager (Java console)" on page 135
- Chapter 8, "Administering the cluster from the command line" on page 235
- Chapter 9, "Configuring resources and applications in VCS" on page 281
- Chapter 10, "Modifying the cluster configuration" on page 369
- Chapter 11, "Predicting VCS behavior using VCS Simulator" on page 405

Chapter 4

About the VCS user privilege model

- About VCS user privileges and roles
- How administrators assign roles to users
- User privileges for OS user groups in secure clusters
- About VCS privileges for users with multiple roles

About VCS user privileges and roles

Cluster operations are enabled or restricted depending on the privileges with which you log on. VCS has three privilege levels: Administrator, Operator, and Guest. VCS provides some predefined user roles; each role has specific privilege levels. For example, the role Guest has the fewest privileges, Cluster Administrator the most.

See “VCS user privileges—administration matrices” on page 685.

About VCS privilege levels

VCS privilege levels include:

- Administrators— Can perform all operations, including configuration options on the cluster, service groups, systems, resources, and users.
- Operators—Can perform specific operations on a cluster or a service group.
- Guests—Can view specified objects.

About user roles in VCS

Table 4-1 lists the predefined VCS user roles, with a summary of their associated privileges.

Table 4-1 User roles in VCS

User Role	Privileges
Cluster Administrator	Cluster Administrators are assigned full privileges, including making configuration read-write, creating and deleting groups, setting group dependencies, adding and deleting systems, and adding, modifying, and deleting users. All group and resource operations are allowed. Users with Cluster Administrator privileges can also change other users’ privileges and passwords. **Note:** Cluster Administrators can change their own and other users’ passwords only after changing the configuration to read/write mode. Cluster Administrators can create and delete resource types.

Table 4-1 User roles in VCS

User Role	Privileges
Cluster Operator	Cluster Operators can perform all cluster-, group-, and resource-level operations, including modifying the user's own password and bringing service groups online. **Note:** Cluster Operators can change their own passwords only if configuration is in read/write mode. Cluster Administrators can change the configuration to the read/write mode. Users with this role can be assigned Group Administrator privileges for specific service groups.
Group Administrator	Group Administrators can perform all service group operations on specific groups, such as bringing groups and resources online, taking them offline, and creating or deleting resources. Additionally, users can establish resource dependencies and freeze or unfreeze service groups. Note that Group Administrators cannot create or delete service groups.
Group Operator	Group Operators can bring service groups and resources online and take them offline. Users can also temporarily freeze or unfreeze service groups.
Cluster Guest	Cluster Guests have read-only access to the cluster, meaning they can view the configuration, but cannot change it. They can modify their own passwords only if the configuration is in read/write mode. They cannot add or update users. Additionally, users with this privilege can be assigned Group Administrator or Group Operator privileges for specific service groups. **Note:** By default, newly created users are assigned Cluster Guest permissions.
Group Guest	Group Guests have read-only access to the service group, meaning they can view the configuration, but cannot change it. The Group Guest role is available in secure clusters only.

About the hierarchy in VCS roles

The following illustration shows the roles and how they overlap with one another.

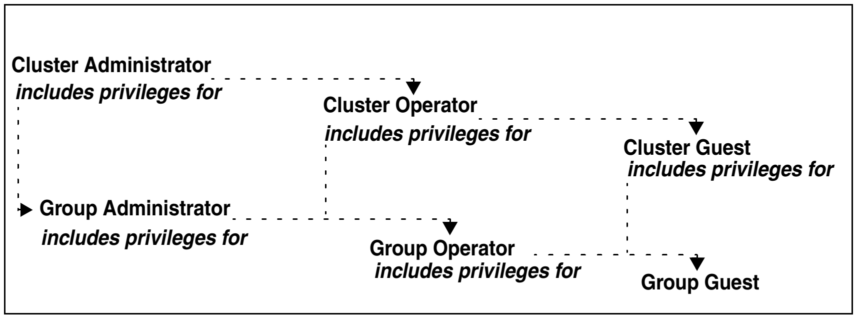

For example, Cluster Administrator includes privileges for Group Administrator, which includes privileges for Group Operator

User privileges for CLI commands

The following concepts apply to users executing commands from the command line:

- Users logged with administrative or root privileges are granted privileges that exceed those of Cluster Administrator, such as the ability to start and stop a cluster.
- If you do not have root privileges, VCS prompts for your VCS user name and password when you execute *haxxx* commands.

User privileges in global clusters

VCS permits a cross-cluster online or offline operation only if the user initiating the operation has one of the following privileges:

- Group Administrator or Group Operator privileges for the group on the remote cluster
- Cluster Administrator or Cluster Operator privileges on the remote cluster

VCS permits a cross-cluster switch operation only if the user initiating the operation has the following privileges:

- Group Administrator or Group Operator privileges for the group on both clusters
- Cluster Administrator or Cluster Operator privileges on both clusters

User privileges in secure clusters

In secure mode, VCS assigns Guest privileges to all native users.

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 using VCS when VCS is running in secure mode.

How administrators assign roles to users

To assign a role to a user, an administrator performs the following tasks:

- Add a user to the cluster, if the cluster is not running in secure mode.
- Assign a role to the user.
- Assign the user a set of objects appropriate for the role. In secure clusters, you can also add a role to an operating system user group.
  See "User privileges for OS user groups in secure clusters" on page 68.

For example, an administrator may assign a user the Group Administrator role for specific service groups. Now, the user has privileges to perform operations on the specific service groups.

You can manage users and their privileges from the command line or from the graphical user interface.

See "Managing VCS users from the command line" on page 249

See "Administering user profiles" on page 173.

User privileges for OS user groups in secure clusters

In secure clusters, you can assign privileges to native users individually or at an operating system (OS) user group level.

For example, you may decide that all users that are part of the OS Administrators group get administrative privileges to the cluster or to a specific service group. Assigning a VCS role to a user group assigns the same VCS privileges to all members of the user group, unless you specifically exclude individual users from those privileges.

When you add a user to an OS user group, the user inherits VCS privileges assigned to the user group.

Assigning VCS privileges to an OS user group involves adding the user group in one (or more) of the following attributes:

- AdministratorGroups—for a cluster or for a service group.
- OperatorGroups—for a cluster or for a service group.
- Guests—for a cluster or for a service group.

For example, user Tom belongs to an OS user group: OSUserGroup1. You can assign VCS privileges to user Tom in the following ways:

To assign privileges	At an individual level, configure attribute	To the OS user group, configure attribute
Cluster Administrator	cluster (Administrators = {tom@domain})	cluster (AdministratorGroups = {OSUserGroup1@domain})
Cluster Operator	cluster (Operators = {tom@domain})	cluster (OperatorGroups = {OSUserGroup1@domain})
Cluster Guest	Cluster *group_name* (Guests = {tom@domain})	cluster *group_name* (Guests = {OSUserGroup1@domain})
Group Administrator	group *group_name* (Administrators = {tom@domain})	group *group_name* (AdministratorGroups = {OSUserGroup1@domain})
Group Operator	group *group_name* (Operators = {tom@domain})	group *group_name* (OperatorGroups = {OSUserGroup1@domain})
Group Guest		group *group_name* (Guests = {OSUserGroup1@domain})

About VCS privileges for users with multiple roles

Table 4-2 describes how VCS assigns privileges to users with multiple roles. The scenarios describe user Tom who is part of two OS user groups: OSUserGroup1 and OSUserGroup2.

Table 4-2 VCS privileges for users with multiple roles

Situation and rule	Roles assigned in the VCS configuration	Privileges that VCS grants Tom
Situation: Multiple roles at an individual level. **Rule:** VCS grants highest privileges (or a union of all the privileges) to the user.	Tom: Cluster Administrator Tom: Group Operator	Cluster Administrator.
Situation: Roles at an individual and OS user group level (secure clusters only). **Rule:** VCS gives precedence to the role granted at the individual level.	Tom: Group Operator OSUserGroup1: Cluster Administrator	Group Operator
Situation: Different roles for different OS user groups (secure clusters only). **Rule:** VCS grants the highest privilege (or a union of all privileges of all user groups) to the user.	OSUserGroup1: Cluster Administrators OSUserGroup2: Cluster Operators	Cluster Administrator
Situation: Roles at an individual and OS user group level (secure clusters only). **Rule:** VCS gives precedence to the role granted at the individual level. You can use this behavior to exclude specific users from inheriting VCS privileges assigned to their OS user groups.	OSUserGroup1: Cluster Administrators OSUserGroup2: Cluster Operators Tom: Group Operator	Group Operator

Getting started with VCS

This chapter provides information on setting up a cluster after installing VCS. For instructions on installing VCS, refer to one of the following documents:

- If you purchased VCS for Network Appliance SnapMirror, refer to one of the following guides depending on the application:
  Veritas Cluster Server Implementation Guide for Microsoft Exchange with NetApp SnapMirror
  or
  Veritas Cluster Server Implementation Guide for Microsoft Exchange Server 2007 with NetApp SnapMirror
  or
  Veritas Cluster Server Implementation Guide for Microsoft SQL with NetApp SnapMirror
  or
  Veritas Cluster Server Implementation Guide for Oracle with NetApp SnapMirror
- If you purchased Veritas Storage Foundation for Windows High Availability, refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

The VCS Configuration wizard configures cluster components including LLT and GAB, the user account for the VCS Helper service, and the Symantec Product Authentication Service. 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 inter-cluster communication for global clusters.

The wizard also modifies and removes cluster configurations.

Prerequisites

- Verify VCS is installed on all systems. Note that VCS must be installed on servers in a Windows Server 2003 domain.
- Verify that each system can access the shared storage devices.
- Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Symantec recommends disabling TCP/IP from private NICs to lower system overhead.
- Verify the public adapters on each node use 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 you must belong to the Domain 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.

Configuring cluster components using the VCS configuration wizard

This section provides instructions for setting up a new cluster. To modify an existing cluster configuration, see "Modifying the cluster configuration" on page 371.

To configure a VCS cluster

1. Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > 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.

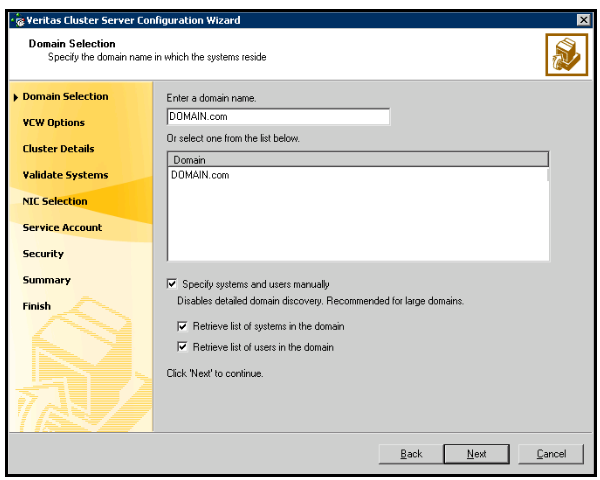

To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 7 on page 75.

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 75. 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.

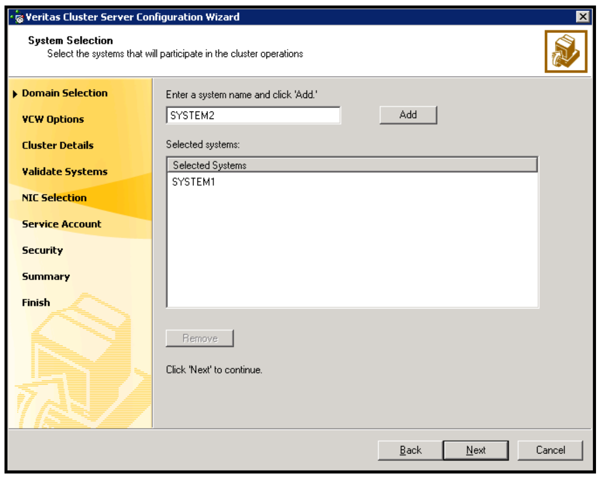

Proceed to step 7 on page 75.

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.

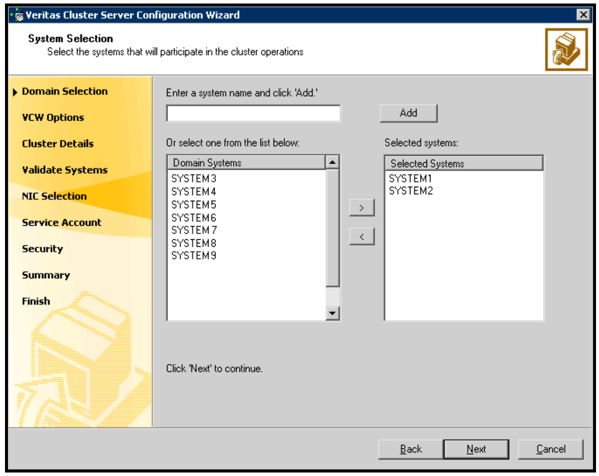

Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

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

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

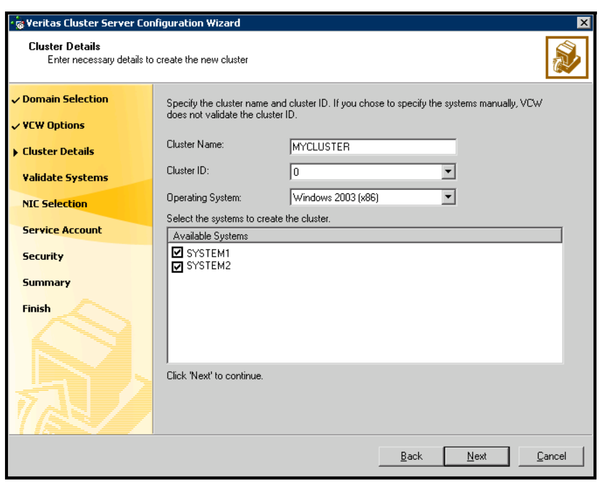

Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

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

Operating System	From the drop-down list, select the operating system 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.

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

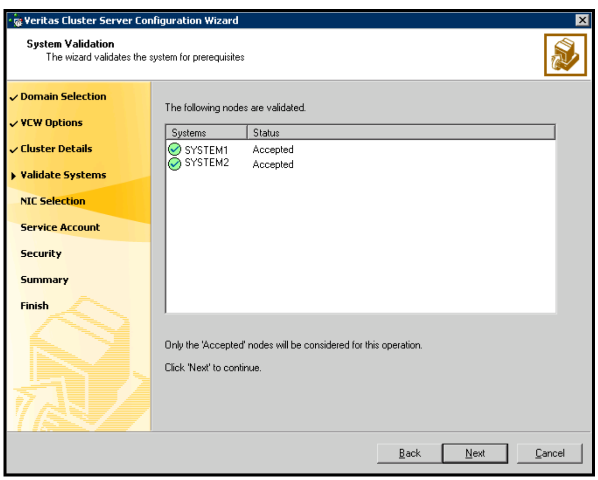

If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in step 8 on page 75, proceed to the next step. Otherwise, proceed to step 11 on page 77.

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

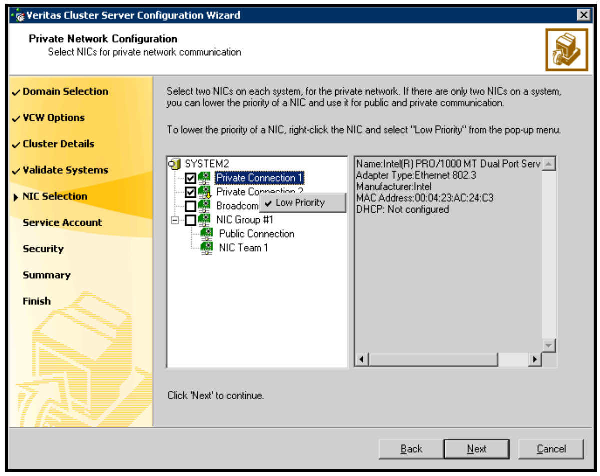

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used for 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.

11 On the VCS Helper Service User Account panel, specify the name of a domain user context 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. Do not use the Administrator account for the VCS Helper service.

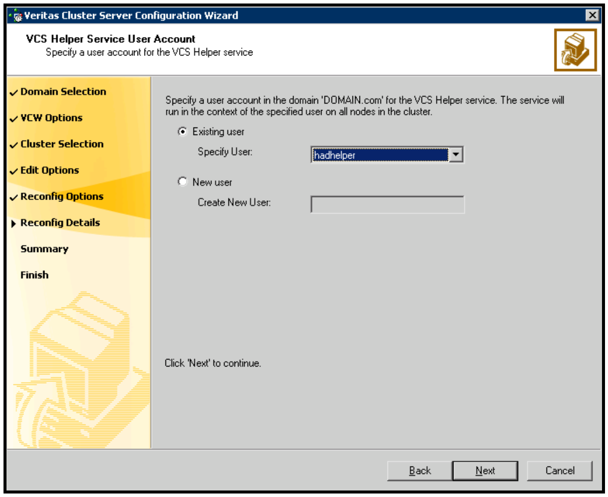

- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list
  - If you chose not to retrieve the list of users in step 4 on page 73, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, 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**.

12 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

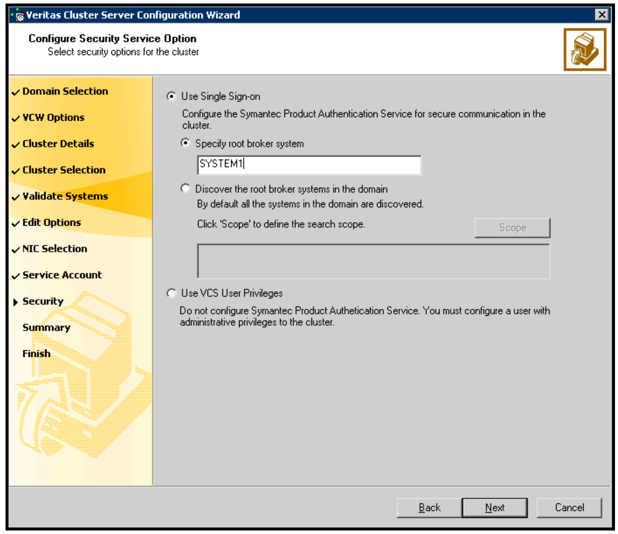

- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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**.
- 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

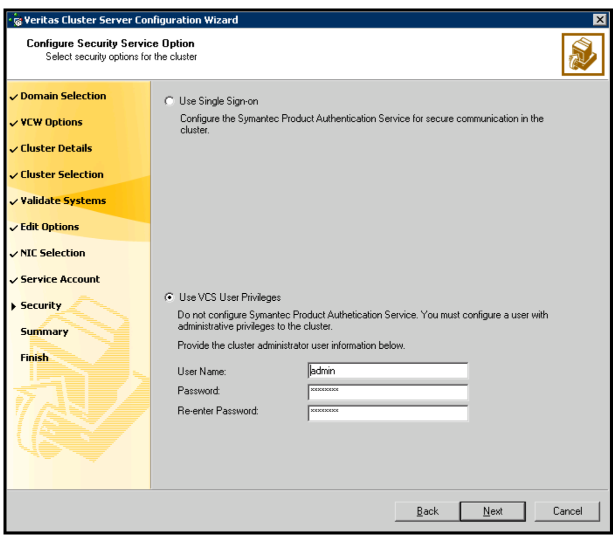

- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

The default user name for the VCS administrator is admin and the default password is password. Both are case-sensitive. 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**.

13 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.

14 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.

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

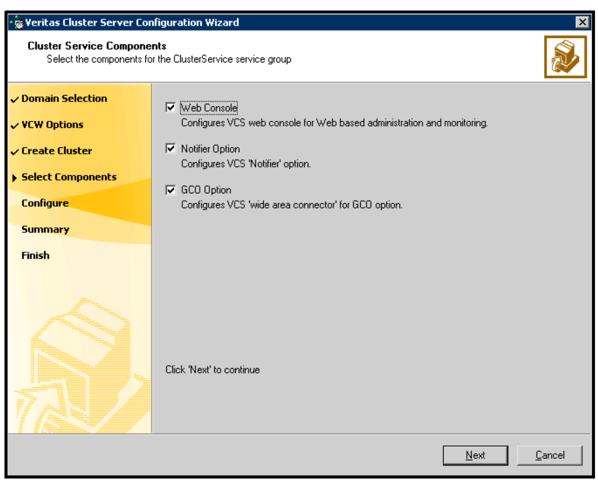

- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
- Check the **GCO Option** check box to configure 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**.

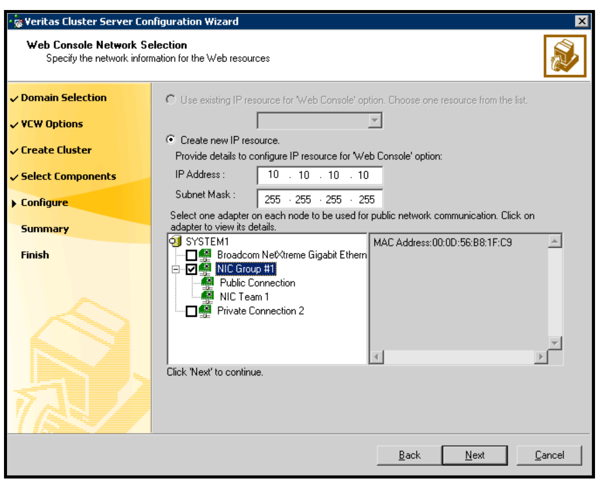

- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
- If you choose to configure a new IP address, type the IP address and associated subnet mask.

- 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 83.
If you chose to configure global cluster components, proceed to "Configuring the wide-area connector process for global clusters" on page 87.
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**.

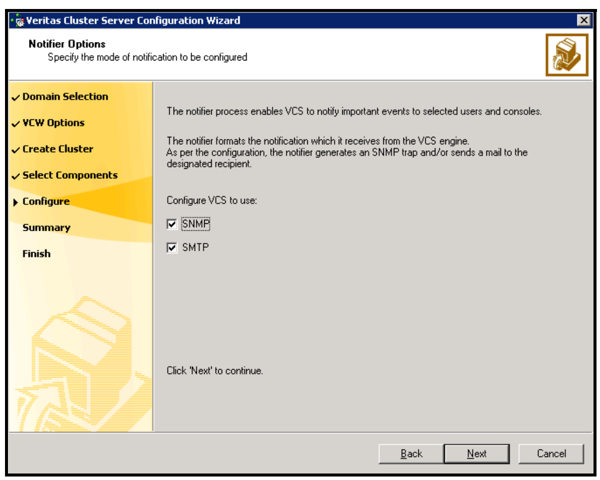

You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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

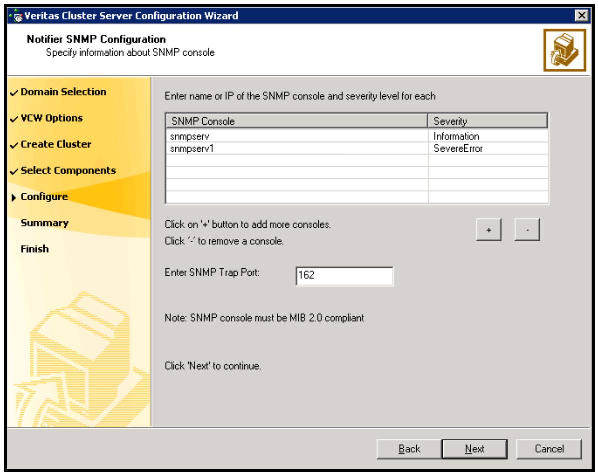

- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click **+** to add a field; click **-** to remove a field.
- 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**.

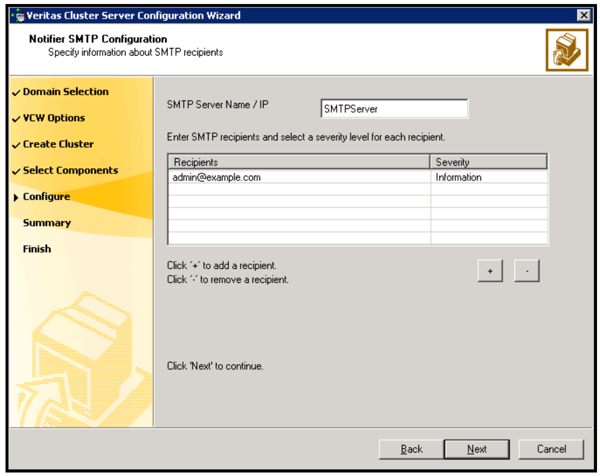

- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity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

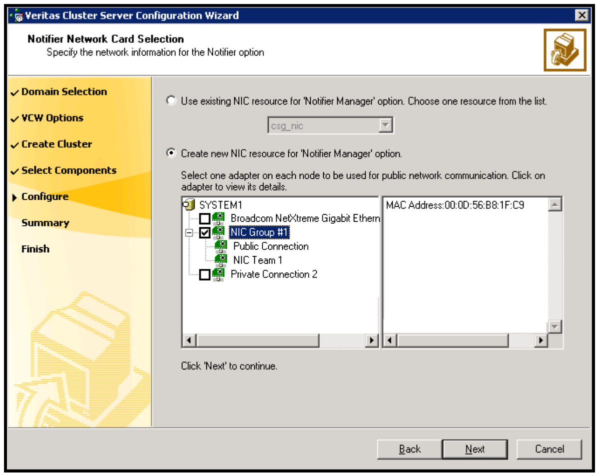

- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists 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 87. 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**.

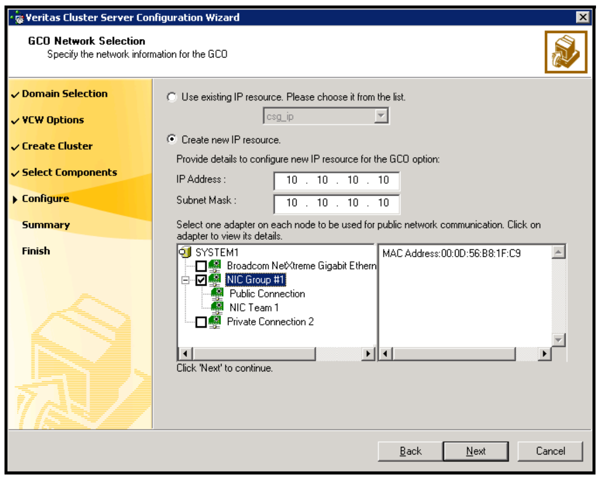

- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
- If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- 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.

Note: The wizard does not set up a global cluster environment; it configures the wide-area connector resource required for inter-cluster communication. See "Setting up a global cluster" on page 553 for more information.

Configuring a cluster from the command line

This section provides instructions for configuring a cluster using the command line. VCS provides a silent configuration utility, `VCWsilent.exe`, which enables you to configure *only* a new cluster. Note that the utility must be used to configure only one cluster at a time.

Preparing for silent configuration

The silent configuration utility requires an XML configuration file, which contains the details of the cluster to be configured. The XML file must be of the following format:

For a non-secure cluster

```
<Operation Type="New">
        <Domain Name="domain_name">
            <SystemList>
                <System Name="sys_name1"/>
                <System Name="sys_name2"/>
                ....
                ....
            </SystemList>
            <Cluster Name="clus_name" ID="clus_ID">
                <Node Name="sys_name1">
                    <LLTLink Name="adp_name_1" MAC="MAC_address_1"
LowPri="pri"/>
                    <LLTLink Name="adp_name_2" MAC="MAC_address_2"
LowPri="pri"/>
                </Node>
                <Node Name="sys_name2">
                    <LLTLink Name="adp_name_1" MAC="MAC_address_1"
LowPri="pri"/>
                    <LLTLink Name="adp_name_2" MAC="MAC_address_2"
LowPri="pri"/>
                </Node>
                ....
                ....
                <Security Type="Non-Secured">
                    <Admin User="admin_user_name"
Password="password"/>
                </Security>
                <HadHelperUser Name="HAD_user_name"
Password="password"/>
            </Cluster>
        </Domain>
</Operation>
```

For a secure cluster

```
<Operation Type="New">
        <Domain Name="domain_name">
            <SystemList>
                <System Name="sys_name_1"/>
                <System Name="sys_name_2"/>
                ....
                ....
            </SystemList>
            <Cluster Name="clus_name" ID="clus_ID">
                <Node Name="node_name_1">
                    <LLTLink Name="adp_name_1" MAC="MAC_address_1"
LowPri="pri"/>
                    <LLTLink Name="adp_name_2" MAC="MAC_address_2"
LowPri="pri"/>
                </Node>
                <Node Name="node_name_2">
                    <LLTLink Name="adp_name_1" MAC="MAC_address_1"
LowPri="pri"/>
                    <LLTLink Name="adp_name_2" MAC="MAC_address_2"
LowPri="pri"/>
                </Node>
                ....
                ....
                <Security Type="Secured">
                    <VxSSRoot Name="root_name"/>
                </Security>
                <HadHelperUser Name="HAD_user_name"
Password="password"/>
            </Cluster>
        </Domain>
</Operation>
```

Copy the relevant format to any text editor and save it with a .xml extension. Replace the variables, shown in italics, with appropriate values. See "Values for element attributes" on page 91 for information about the variables and their possible values.

You may also refer to a sample XML file presented at "Sample XML configuration: Two node secure cluster" on page 93.

Values for element attributes

The following table describes the variables used in the XML format and their possible values:

Table 5-3 VCWsilent - variables and values

Variables	Description
domain_name	Replace this with the fully qualified name of a domain in which the systems reside.
sys_name_<n>*	Replace this with name of the system in the domain for which relevant information will be discovered. **Note:** For each system, you must have a *System* child element under the *SystemList* element.
clus_name	Replace this with the name of the cluster to be created.
clus_ID	Replace this with the cluster ID. Make sure you specify a unique cluster ID between 0 and 255.
node_name_<n>*	Replace this with the name of the system that will be part of the cluster. Make sure you provide system names from the list of systems that are specified under the *SystemList* element. For example, if you specified SysA and SysB in the *SystemList* element, you may specify one or both the systems for the node names. However, you should not specify another system, say SysC, which was not specified in the *SystemList* element. **Note:** For each node, you must have a *Node* child element along with the *LLTLink* subchild element under the *Cluster* element.
adp_name_<n>*	Replace this with the name of the adapter on which the LLT link will be configured. **Note:** For each node, you must specify a minimum of two adapters. Each adapter must be specified as an attribute of the *LLTLink* element.
MAC_address_<n>*	Replace this with the MAC address of the adapter.
Pri	Replace this with either "1" or "0". Value "1" indicates the adapter is assigned a low priority. Value "0" indicates otherwise. You may assign a low priority to an adapter to use it for both private and public network communication.

Table 5-3 VCWsilent - variables and values (continued)

Variables	Description
admin_user_name	Replace this with a user name for the cluster administrator. You may use this user name to log on to a cluster using Cluster Manager. **Note:** This is applicable only for a non-secure cluster.
root_name	Replace this with a system name that will be configured as the root broker for all the systems in the cluster. If you provide a system that is not specified under the *Node* element, make sure the system is configured as a root broker. **Note:** This is applicable only for a secure cluster.
HAD_user_name	Replace this with a domain user name in whose context the VCS Helper service will run. The VCS High Availability Daemon, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.
password	Replace this with an encrypted password. See "Encrypting passwords" on page 92 for instructions.

* "n" is the sequence number for the systems, nodes, adapaters, and MAC addresses.

Encrypting passwords

Before specifying passwords in the XML configuration file, you must encrypt them using the `vcsencrypt` utility.

To encrypt a password

Perform these steps for all the passwords to be specified in the XML file.

1. Run the `vcsencrypt` utility. Type the following on the command line.
   ```
   C:\> vcsencrypt -agent
   ```
2. The utility prompts you to enter the password twice. Enter the password and press Return.
   ```
   Enter New Password:
   Enter Again:
   ```
3. The utility encrypts the password and displays the encrypted password.
4. Specify this encrypted password in the XML file.
   Copy the encrypted password for future reference.

Sample XML configuration: Two node secure cluster

This configuration file can be used to create a two node secure cluster with systems SYSTEM11 and SYSTEM2.

```
<Operation Type="New">
    <Domain Name="DOMAIN.com">
        <SystemList>
            <System Name="SYSTEM1"/>
            <System Name="SYSTEM2"/>
        </SystemList>
        <Cluster Name="MYCLUSTER" ID="0">
            <Node Name="SYSTEM1">
                <LLTLink Name="Adapter0" MAC="00:03:47:08:91:56"
LowPri="0"/>
                <LLTLink Name="Adapter1" MAC="00:03:47:08:91:C6"
LowPri="0"/>
            </Node>
            <Node Name="SYSTEM2">
                <LLTLink Name="Adapter0" MAC="00:03:47:08:91:CC"
LowPri="0"/>
                <LLTLink Name="Adapter1" MAC="00:03:47:08:94:4E"
LowPri="0"/>
            </Node>
            <Security Type="Secured">
                <VxSSRoot Name="SYSTEM1"/>
            </Security>
            <HadHelperUser Name="Administrator" Password="hvnTkvK"/>
        </Cluster>
    </Domain>
</Operation>
```

Performing a silent configuration

Before running the silent configuration utility, `VCWsilent.exe`, make sure you meet the prerequisites listed under "Prerequisites" on page 72.

To run the silent configuration utility

You may run the utility from any system in the domain, irrespective of whether the system will be part of the cluster being configured.

1. From the command line, navigate to the directory containing the XML configuration file and run the `VCWsilent` utility. Type the following on the command line:

   ```
   C:\<XML_file_location>> VCWsilent <name of XML file>
   ```

   To view the progress of the silent configuration, use the "-v" option. Type the following on the command line:

   ```
   C:\<XML_file_location>> VCWsilent <name of XML file> -v
   ```

2. If the cluster is successfully configured, the following message is displayed:

   ```
   Silent configuration was successful.
   ```

   If the silent configuration fails, an error message is displayed. Review the message associated with the failure and rerun the utility after rectifying the problem.

Analyzing a domain

The VCS Configuration Wizard analyzes a domain and creates a report about clusters and systems in the domain. The wizard retrieves information about whether VCS is installed and configured on systems, and about the network configuration of cluster nodes.

To analyze a domain and create an analysis report

1. Start the VCS Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
2. Read the information on the Welcome panel and click **Next**.
3. On the Configuration Options panel, click **Analyze Domain** and click **Next**.
4. On the Domain Selection panel, select or type the name of the domain to be analyzed and click **Next**. Review the information in the message box and click **Yes**.
5. On the Domain Analysis for VCS panel, select the clusters and systems to view their respective details.

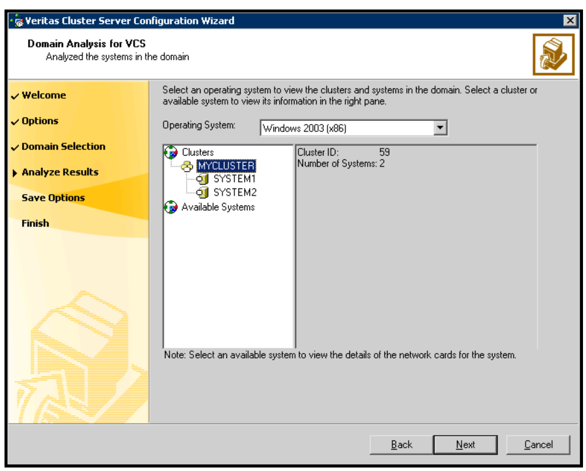

- Select an operating system.
- Double-click **Clusters** to view the clusters running the selected operating system.
- Select a cluster name to view the cluster ID and the number of nodes in the cluster.
- Double-click a cluster name to view the nodes in the cluster.

- Select a node to view its node ID and the number of LLT links configured on the node.
- Double-click **Available Systems** to view the systems that have VCS installed but not configured. You can use the wizard to add these nodes to a new cluster, or to form a new cluster with these nodes.
- Select a system to view the install path for VCS and the number of adapters on the system.
- Click **Next**.

6 Specify the options for generating a domain analysis report.

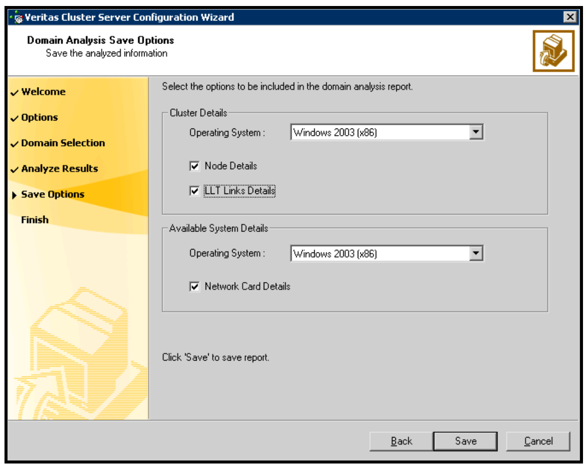

- In the Cluster Details box, select the operating system that the cluster nodes are running.
- Check the **Node Details** check box to obtain information about the nodes in the cluster.
- Check the **LLT Link Details** check box to obtain information about the LLT links between the cluster nodes.
- In the Available System Details box, select the operating system that the systems are running.
- Check the **Network Card Details** check box to obtain information about the network adapters on the systems.
- Click **Save**.

7 The wizard generates the report in the XML format and saves it under the user profile in whose context the wizard was run, typically `C:\Documents and Settings\`*`<user profile>`*`\Application Data\VERITAS\Cluster Server\VCWAnalysis.xml`.

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

Chapter 6

This chapter includes the following topics:

- Administering a cluster
- Administering service groups in a cluster
- Administering resources in a cluster
- Administering resource types in a cluster
- Administering systems in a cluster
- Administering attributes in a cluster
- Conducting a search

Administering a cluster

The VCS Management Console enables you to initially set and then modify the parameters of the VCS cluster configuration. After opening the configuration, you can save it to disk. Use the Cluster:Summary view to administer the cluster.

To navigate to the Cluster:Summary view

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Summary**.
In the Cluster:Summary view, you can choose a task from the task pane or select another cluster-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the cluster.

Opening the configuration

You can modify a read-only configuration file to a read-write file by opening the configuration. This task is available only when the cluster configuration has been saved in read-only mode.

To open a configuration

1 In the Cluster:Summary view, in the Configuration task panel, click **Open Configuration**.

2 In the Open Configuration dialog box, click **OK** to confirm that you want to open the configuration in read-write mode.

Saving the configuration

After updating the VCS cluster configuration, save it to disk in either read-only or read-write mode.

To save the configuration

1 In the Cluster:Summary view, in the Configuration task panel, click **Save Configuration**.

2 In the Save Configuration dialog box, check the **Make Configuration Read-Only** option to save the configuration in read-only mode, or leave it clear to save the configuration in read-write mode.

3 Click **OK**.

Modifying cluster attributes

Edit the attributes of a cluster. You must have the role of cluster administrator to edit cluster attributes.

To edit a cluster attribute using the attribute table

1. In the Cluster:Summary view, on the secondary tab bar, click **Attributes**.
2. In the All attributes for cluster table, locate the line item for the attribute you want to edit.
3. In the Edit column, click the**...** (edit) button.
4. In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:
   - The attribute value
     In the Values field, either click a value to change the it, or click the **+** (plus) button to add a value. Click a value and then click the **-** (minus) button to remove a value.
     You can enter more than one value for some attributes. Use the up- and down-arrow buttons to scroll among multiple values.
   - The systems on which to change the attribute
     Specify a setting for the **Apply value to** option. **Apply to all nodes** applies the changes to all systems listed the System List, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the System List.

Configuring the cluster connector for secure-mode clusters

You can use the cluster connector on a cluster that is not configured in secure mode. If you later decide to configure that cluster in secure mode, you must also configure the cluster connector to use only secure credentials to communicate with HAD, the VCS engine.

The script enable_security.pl enables you to configure the cluster connector to use secure-mode communication. This script is installed on a cluster along with the cluster connector and is available on any cluster node in the following path:

```
C:\Program Files\VERITAS\Cluster Management Console\
ClusterConnector\bin
```

Note: You must run the enable_security.pl script to configure the cluster connector into secure mode. However, you are not required to run any additional scripts (for cluster connector) if you revert the cluster back out of secure mode.

To configure cluster connector to use secure-mode communications

1 On a cluster node, obtain a command prompt and enter the following command:

```
enable_security.pl password
```

Replace `password` with the password that you want to use to create the CMC_HAD@CMC_SERVICES account.
See "Internal user accounts" on page 31.

2 Repeat step 1 on each node of the cluster and then restart the CMC service group.

Administering service groups in a cluster

The Cluster Management Console enables you to add and configure a service group according to the requirements of the application that the service group supports. Use the following views to administer a service group:

- Cluster:Groups
  This is a cluster-level view of all service groups in the cluster.
- Group:Summary
  This is the first view at the service group level in the VCS Management Console.

To navigate to the Cluster:Groups view

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Summary**, and then click **Groups**.
In the Cluster:Groups view, the Service Groups Listing table lists all service groups in the selected cluster. The task panels contain several tasks that you can perform in a one-to-many fashion on the listed service groups.

To navigate to the Group:Summary view

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Summary** and then click **Groups**.

3 In the Cluster:Groups view, in the Groups Listing table, click a linked service group name.
In the Group:Summary view, you can choose a task from the task pane or select another service-group-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering a single service group.

Adding a service group to a cluster

Add and configure a service group as a member of the cluster. You must have the role of cluster administrator to add a service group to a cluster.

To add a service group to a cluster

1. In the Cluster:Groups view, in the Configuration task panel, select **Add Service Group**.
2. In the Service Group Configuration wizard, read the introductory text and then click **Next**.
3. In the **Service Group Configuration** panel, specify the following details of the new service group and then click **Next**:
   - The service group name
     Enter the name into the **Service Group Name** text box. Service group names must:
     - Consist of only one string
     - Not contain special characters like . (period) or & (ampersand).
     - Begin with an alphabetic character.
       The remaining characters in the name can be letters, digits, hyphens, or underscores.
   - The service group type
     Click either **Failover** or **Parallel** to specify the service group type.
     See “Types of service groups” in the Veritas Cluster Server user’s guide.
4. In the System List Configuration panel, specify the following details for the target system and then click **Next**. (The target system is the system on which you configure the service group.)
   - The target system name
     Find the system you want in the list under Select and Add Systems. To quickly find a specific system, click **Filter**, type a few letters of the system name, and then click **Apply**.
     To select a system, click the system name under Select and Add Systems and then click the right-arrow button to move it under Selected Systems. You can select multiple systems. To reject a system, click the system name under Selected Systems and then click the left-arrrow button. This action moves the system back under Select and Add Systems.
   - The startup system
     Under Select and Add Systems, check the **Startup** option next to a system name if you want the service group to startup automatically on that system.

5 In the Resource Creation dialog box, specify the following details about the resources and the service group and then click **Next**:

- The resource name
  Enter the name of the resource that you want to add to the service group into the Resource Name text box. Resource names must:
  - Consist of only one string
  - Not contain special characters like . (period) or & (ampersand).
  - Begin with an alphabetic character.
    The remaining characters in the name can be letters, digits, hyphens, or underscores.
- The resource type
  In the Resource Type drop-down list, select the type of the new resource.
- The resource startup state
  If you want the resource to start in an enabled state after the service group comes online, check **Enable Resource**.
- The resource list for the service group
  Click **Add Resource** to add the resource to the member resource list for the service group. The resource list is used to populate the Resource List table.
  In the Resource List table, you can edit resource attributes or delete a resource using the buttons in the Edit and Delete columns.

6 If you need to link one or more resources, provide the following relationship details in the Link/Unlink Resource panel and then click **Finish**. Otherwise, click **Finish** now to exit the wizard and create the new service group.

- The parent resource
  In the Parent Resource drop-down list, specify which resource must be brought online after the child resource.
- The child resource
  In the Child Resource list box, select which resource should be brought online first, and then click **Add Link**.

Deleting a service group from a cluster

Manually remove a service group from cluster membership. You must have the role of cluster administrator to remove a service group from a cluster.

You cannot delete a service group if it is a child in a dependency. You must first remove the dependency and then delete the service group.

See “Unlinking service groups” on page 113.

To delete a service group from a cluster

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to delete.
2. In the Group:Summary view, in the Configuration task panel, click **Delete Service Group**.
3. In the Delete Service Group dialog box, click **OK** to confirm that you want to delete the specified service group.

Modifying a service group

Edit a service group configuration. You cannot edit the service group name or type using this task. To change the name or type, you must first remove the service group and then add it back using the desired name and type. You must have the role of cluster administrator or service group administrator to edit a service group.

To edit a service group configuration

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to modify.
2. In the Group:Summary view, in the Configuration task panel, click **Modify Service Group**.
3. In the Service Group Configuration wizard, read the introductory text and then click **Next**.
4. In the Service Group Configuration panel, the service group name and type appear but cannot be modified using this task. Click **Next**.
5. In the System List Configuration dialog box, specify the following details for the target systems and then click **Next**. (The target systems are the systems on that can host the service group.)
   - The target system names
     Find the system you want in the list under Select and Add Systems. To quickly find a specific system, click **Filter**, type a few letters of the system name, and then click **Apply**.
     To select a system, click the system name under Select and Add Systems and then click the right-arrow button to move it under Selected Systems. You can select one or more systems.
     To reject a system, click the system name under Selected Systems and then click the left-arrrow button to move it back under Select and Add Systems.

- The startup system

  Under Selected Systems, check the **Startup** option next to a system name if you want the service group to startup automatically on that system.

- The system order of preference

  Use the up- and down-arrow buttons to the right of the Selected Systems table to set the order of systems in the list. To change the position of a system in the list, click the system name in the Selected Systems table and then click the up- or down-arrow button to move it up or down in the list.

  VCS attempts to bring the service group online on the first system in the system list, unless you have selected a different system as the startup system using the Startup option. If the service group cannot be brought online on the first system or on the designated startup system, VCS does one of the following:

  - If the first system is also the startup system, VCS attempts to bring the service group online on each remaining system in the order listed.
  - If the first system and the startup system are different, VCS attempts to bring the service group online on the first system in the list. If the first system fails, VCS attempts to bring the service group online on each remaining system in the order listed.

6 In the Resource Creation panel, specify the following details about the resources that belong to the service group and then click **Next**:

- The member resource list

  Enter the name of any new resources that you want to add to the service group into the **Resource Name** text box. Resource names must:

  - Consist of only one string
  - Not contain special characters like . (period) or & (ampersand).
  - Begin with an alphabetic character.

    The remaining characters in the name can be letters, digits, hyphens, or underscores.

- The resource startup state

  If you want the resource to start in an enabled state after the service group comes online, check **Enable Resource**. Click **Add Resource** to add the resource to the **Resource List**.

- The resource attribute values or resource removal

  In the Resource List table, you can edit resource attributes or delete a resource using the buttons in the Edit and Delete columns.

  See "Adding a resource to a service group" on page 119.

7 If you need to link one or more resources, provide the following relationship details in the Link/Unlink Resource panel and then click **Finish**. Otherwise, click **Finish** now to exit the wizard.

- The parent resource
  In the Parent Resource drop-down list, specify which resource must be brought online after the child resource.
- The child resource
  In the Child Resource list box, select which resource should be brought online first, and then click **Add Link**.

Modifying the system list of a service group

Edit the list of systems that can host a service group. You must have the role of cluster administrator or service group administrator to modify the system list of a service group.

To modify the system list of a service group

1 In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to modify.

2 In the Group:Summary view, in the Configuration task panel, click **Modify System List**.

3 In the System List Configuration dialog box, specify the following details for the *target* systems and then click **Finish**. (The target systems are the systems on that can host the service group.)

- The target system names
  Find the system you want in the list under Select and Add Systems. To quickly find a specific system, click **Filter**, type a few letters of the system name, and then click **Apply**.
  To select a system, click the system name under Select and Add Systems and then click the right-arrow button to move it under Selected Systems. You can select one or more systems.
  To reject a system, click the system name under Selected Systems and then click the left-arrrow button to move it back under Select and Add Systems.
- The startup system
  Under Selected Systems, check the **Startup** option next to a system name if you want the service group to startup automatically on that system.

- The system order of preference
  Use the up- and down-arrow buttons to the right of the Selected Systems table to set the order of systems in the list. To change the position of a system in the list, click the system name in the Selected Systems table and then click the up- or down-arrow button to move it up or down in the list.
  VCS attempts to bring the service group online on the first system in the system list, unless you have selected a different system as the startup system using the Startup option. If the service group cannot be brought online on the first system or on the designated startup system, VCS does one of the following:
  - If the first system is also the startup system, VCS attempts to bring the service group online on each remaining system in the order listed.
  - If the first system and the startup system are different, VCS attempts to bring the service group online on the first system in the list. If the first system fails, VCS attempts to bring the service group online on each remaining system in the order listed.

Bringing a service group online

Manually put a service group into a responsive, functioning state. You must have the role of cluster operator or service group operator to bring a service group online.

You can bring a service group online on a specific system, or you can bring the service group online "anywhere." If you select the Anywhere option, the service group is brought online on the first available system in the cluster.

To bring a service group online with the opportunity to specify one system

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the offline service group that you want to bring online.
2. In the Group:Summary view, in the Operations task panel, click **Online**.
3. This step is necessary only if the system list for the service group has more than one member system.
   In the Online Service Group dialog box, select the system on which to bring the service group online, or click **Anywhere**.
   The system choices in this step are populated using the system list for the service group. The Anywhere option causes this task to try each system in the list until the service group is successfully brought online.
4. Click **OK** to confirm that you want to bring the service group online.

To bring one or more service groups online anywhere

1. In the Cluster:Groups view, in the Groups Listing table, check the check box preceding the line item for the offline service group that you want to bring online. Select all service groups by checking the check box at the top.
2. In the Operations task panel, click **Online Anywhere**.
3. In the **Online Service Group(s)** dialog box, click **OK** to confirm that you want to bring the selected service groups online on any system in the cluster.
   The candidate systems are members of the system list for each service group. The Online Anywhere task tries each system in a system list until every service group is successfully brought online.

Taking a service group offline

Manually put a service group into an unresponsive, nonfunctioning state. You must have the role of cluster operator or service group operator to take a service group offline.

You can take a service group offline on a specific system, or you can take the service group offline "anywhere". If you select the Anywhere option, the service group is taken offline on the first system in the cluster on which it is found.

To take a service group offline with the opportunity to specify one system

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the online service group that you want to take offline.
2. In the Group:Summary view, in the Operations task panel, click **Offline**.
3. This step is necessary if only if the service group is online on more than one system.
   In the Offline Service Group dialog box, select the system on which to take the service group offline, or click **Anywhere**.
4. Click **OK** to confirm that you want to take the service group offline.

To take one or more service groups offline anywhere

1. In the Cluster:Groups view, in the Clusters Listing table, check the check box preceding the line item for each online service group that you want to take offline. Select *all* service groups by checking the check box at the top.
2. In the Operations task panel, click **Offline Anywhere**.
3. In the Offline Service Group(s) dialog box, click **OK** to confirm that you want to take the selected service groups offline anywhere (on any system) in the cluster.

Switching a service group

Manually switch a service group to another system in the cluster.

To switch a service group

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to switch.
2. In the Group:Summary view, in the Operations task panel, click **Switch**.
3. In the Switch Service Group dialog box, select the destination system for the service group in the drop-down list. Select **Anywhere** to switch the service group to any other system in the cluster.
   The system choices in this step are populated using the system list for the service group. The Anywhere option causes this task to try every other system in the list until the service group is successfully brought online.
4. If insufficient room exists on the target system, you can check the **Evacuate lower-priority service groups** option. This option makes room for the service group being switched by taking lower-priority service groups offline.
5. Click **OK**.

Freezing a service group

Prevent a service group from failing over to another system. Freezing a service group stops all online and offline operations on the service group.

To freeze a service group

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to freeze.
2. In the Group:Summary view, in the Operations task panel, click **Freeze**.
3. In the Freeze Service Group dialog box, click **OK** to confirm that you want to freeze the specified service group. If you want the service group to remain frozen until explicitly unfrozen, check the **Freeze Persistently** option. Otherwise, the service group unfreezes when either the management server or the VCS Management Console is restarted.

To freeze one or more service groups

1. In the **Cluster:Groups** view, in the **Groups Listing** table, check the check box preceding the line item for each unfrozen service group that you want to freeze. Select *all* service groups by checking the check box at the top.
2. In the **Operations** task panel, click **Freeze**.

3 In the **Freeze Service Group(s)** dialog box, click **OK** to confirm that you want to freeze the specified service groups.

Unfreezing a service group

Return a frozen service group to normal failover behavior. Unfreezing a service group restarts online and offline operations on the service group.

To unfreeze a service group

1 In the Cluster:Groups view, in the Groups Listing table, click the linked name of the frozen service group that you want to unfreeze.

2 In the Group:Summary view, in the Operations task panel, click **Unfreeze**.

3 In the Unfreeze Service Group dialog box, click **OK** to confirm that you want to unfreeze the specified service group.

To unfreeze one or more service groups

1 In the Cluster:Groups view, in the Groups Listing table, check the check box preceding the line item for each frozen service group that you want to unfreeze. Select *all* service groups by checking the check box at the top.

2 In the Operations task panel, click **Unfreeze**.

3 In the **Unfreeze Service Group(s)** dialog box, click **OK** to confirm that you want to unfreeze the specified service groups.

Enabling a service group

Enable a disabled service group so that it can be brought online. A service group that is manually disabled during a maintenance procedure must first be enabled before it can be brought online.

To enable a service group

1 In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to enable.

2 In the Group:Summary view, in the Operations task panel, click **Enable**.

3 In the Enable Service Group dialog box, use the drop-down menu to select a system on which to enable the service group. Select **all** to enable the service group on all systems in the cluster.

4 Click **OK**.

Disabling a service group

Disable a service group to prevent it from coming online. Use disablement to temporarily prevent agents from monitoring a service group on a system undergoing maintenance operations.

To disable a service group

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the enabled service group that you want to disable.
2. In the Group:Summary view, in the Operations task panel, click **Disable**.
3. In the Disable Service Group dialog box, use the drop-down menu to select a system on which to disable the service group. Select **all** to disable the service group on all systems in the cluster.
4. Click **OK**.

Linking service groups

Create dependencies among service groups.

See "About service group dependencies" in the *Veritas Cluster Server user's guide*.

To link two service groups with a parent-child dependency

1. In the Cluster:Groups view, in the Configuration task panel, click **Link Service Groups**.
2. In the Link Service Groups dialog box, specify the following details of the dependency and then click **OK**:
   - The service groups that you intend to serve as the parent service group and the child service group
     Use the corresponding drop-down menus for this specification.
   - The relationship type
     In a Local dependency, an instance of the parent service group depends on an instance of the child service group being online or offline on the same system, depending on the category of group dependency. In a Global dependency, an instance of the parent service group depends on one or more instances of the child service group being online on any system.
   - The dependency category
     In a Soft dependency, VCS imposes minimal constraints while bringing the parent and child service groups online and taking them offline. In a Firm dependency, VCS takes the child service group offline before

taking the parent service group offline when the child service group faults. In a Hard dependency, VCS takes the parent service group offline before taking the child service group offline when the child service group faults.
Hard dependencies are designed for use with Veritas Volume Replicator in disaster recovery configurations. In these configurations, the application is in the parent service group and the replication resources are in the child service group.

Viewing service group dependencies

View a graphical representation of dependencies that exist among service groups.

View a graphical representation of dependencies that exist among service groups.

To view service group dependencies and get status information

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Summary** and then click **Group Dependency**.
This view presents a graphical representation of all currently configured service groups and the dependencies that exist among them.

3 Click a service group icon to get information about the service group that it represents.

Unlinking service groups

Remove dependencies among service groups.

To unlink two service groups that share a parent-child dependency

1 In the Cluster:Groups view, in the Configuration task panel, click **Unlink Service Groups**.

2 In the Unlink Service Groups dialog box, select the parent group and the child group from the corresponding drop-down menus and then click **OK**.

Adding systems to a service group

Add systems to the system list of a service group. You can add systems to the system list of a service group when you add or modify the service groups in a cluster. You add or modify service groups using the Service Group Configuration wizard.

See "Adding a service group to a cluster" on page 103.

See "Modifying a service group" on page 105.

Deleting systems from a service group

Delete systems from the system list of a service group. You can delete systems from the system list of a service group when you add or modify the service groups in a cluster. You add or modify service groups using the Service Group Configuration wizard.

See "Adding a service group to a cluster" on page 103.

See "Modifying a service group" on page 105.

Faulting a service group

Manually simulate a fault on a service group to test failover behavior. This operation is available only when using the cluster simulator.

To fault a single service group

1. In the Cluster:Groups view, in the Groups Listing table, click the linked name of the service group that you want to fault.
2. In the Group:Summary view, in the Simulation task panel, click **Fault Service Group**.
3. In the **Fault Service Group** dialog box, select the systems on which you want to fault the service group from the drop-down menu.
   This drop-down menu is populated using the system list of the service group. The system list of a service group contains a list of all systems on which the service group is configured.
4. Click **OK**.

To fault one or more service groups

1. In the Cluster:Groups view, in the Groups Listing table, check the check box preceding the line item for each service group that you want to fault. Select *all* service groups by checking the check box at the top.
2. In the Simulation task panel, click **Fault Service Groups**.
   In the **Fault Service Group(s)** dialog box, use the appropriate drop-down menu to select the systems on which you want to fault each service group. Each drop-down menu is populated using the system list of each service group. The system list of a service group contains a list of all systems on which the service group is configured.
3. Click **OK**.

Clearing a faulted service group

Clear a service group to remove the resource faults within the group. You must have the role of cluster administrator or service group administrator to clear a faulted service group.

This task makes the service group available to be brought online. A resource fault in a service group may occur in several situations, such as a power failure or faulty configuration.

To clear a faulted service group

1. In the Group:Summary view, in the Operations task panel, click **Clear Fault**.
2. In the Clear Faulted Group dialog box, use the drop-down menu to select the system on which to clear the service group. To clear the group on all systems, click **All Systems**.
3. Click **OK**.

To clear one or more service groups

1. In the Cluster:Groups view, in the Groups Listing table, check the check box preceding the line item for each service group that you want to clear. Select *all* service groups by checking the check box at the top.
2. In the Operations task panel, click **Clear Fault**.
   In the Clear Service Group(s) Faults dialog box, use the appropriate drop-down menu to select the systems on which you want to clear each service group. Each drop-down menu is populated using the system list of each service group. The system list of a service group contains a list of all systems on which the service group is configured.
3. Click **OK**.

Modifying service group attributes

Edit the attributes of a service group. You must have the role of cluster administrator or service group administrator to edit attributes.

To edit a service group attribute using the attribute table

1. In the Group:Summary view, click **Attributes**.
2. In the All attributes for group table, locate the line item for the attribute you want to edit.
3. In the Edit column, click the **...** (edit) button.
4. In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:

- The attribute value
  In the Values field, either click a value to change the it, or click the **+** (plus) button to add a value. Click a value and then click the **-** (minus) button to remove a value.
  You can enter more than one value for some attributes. Use the up- and down-arrow buttons to scroll among multiple values.
- The systems on which to change the attribute
  Specify a setting for the **Apply value to** option. **Apply to all nodes** applies the changes to all systems listed the **System List**, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the **System List**.

To edit a service group attribute using the task pane

1. In the Cluster:Summary view, in the Groups Listing table, check the check box preceding the line item for each service group with the attribute that you want to edit. Select all service groups by checking the check box at the top.
2. In the Configuration task panel, click **Edit Attribute**.
3. Specify the attribute details.
   See "Modifying service group attributes" on page 115.

Flushing a service group

Postpone the resolution of resource-related problems that occur while a service group is being brought online or taken offline. You must have the role of cluster administrator or service group administrator, or service group operator to flush a service group.

Flushing a service group halts the current operation on the resources yet to be brought online or taken offline. Flushing a service group typically leaves the cluster in a partial state. After flushing a service group, you can resolve the issue with the problematic resource.

To flush a service group

1. In the Group:Summary view, in the Operations task panel, click **Flush**.
2. In the Flush Service Group dialog box, select the system on which you want to flush the service group from the drop-down menu.
   The menu is populated with system from the service group system list, which contains the systems on which the service group is configured.
3. Click **OK**.

Adding a RemoteGroup resource

A RemoteGroup resource enables you to monitor and manage a service group on a remote cluster. You must have the role of cluster administrator or service group administrator to configure a remote group resource.

Use the RemoteGroup agent to create a dependency between a local and remote service group. Typically, the RemoteGroup resource depends on an IP resource in the local service group.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on this agent.

See "Configuring the RemoteGroup agent" in the *Veritas Cluster Server Administrator's Guide*.

Note: Although this topic refers to a remote service group, the Add Remote Group task is not part of VCS global cluster configuration. This task is intended to provide a simple way to monitor and manage a service group on another cluster.

To configure a remote group resource on a local cluster

1. In the Group:Summary view, in the Configuration task panel, click **Add Remote Group Resource**.
2. In the **Remote Group Resource** wizard, read the introductory information and then click **Next**.
3. In the RemoteGroup Resource dialog box, enter a name in the **RemoteGroup Resource Name** box and then click **Next**.
4. In the Remote Cluster Authentication dialog box, enter the host IP address and all of the required user authentication information. If you are logging in to a secure cluster, check **Select to enter domain name for Secure Cluster** and then enter the domain name in the box provided.
5. In the Create RemoteGroup Resource dialog box, specify the following monitoring details and then click **Finish**:
   - The remote service group
     In the Remote Service Group drop-down list, select the remote service group. This list is populated using the service groups configured in the remote cluster that you specified.
   - The type of monitoring you want to perform.
     The following options are available on the **Type of Monitoring** drop-down menu:

- OnlineOnly
  The RemoteGroup resource only brings the remote service group online. The RemoteGroup resource cannot take the remote service group offline.
- MonitorOnly
  The RemoteGroup resource only monitors the state of the remote service group. The RemoteGroup resource cannot online or offline the remote service group.
- OnOff
  The RemoteGroup resource brings the remote service group online or takes it offline.

- Monitoring details
  The System Based Monitoring Details setting has two options, Any and One-to-One Mapping. See the *Veritas Cluster Server Bundled Agents Reference Guide* for information on this setting.

After you configure the RemoteGroup resource, you can view the service group dependency in the Cluster:Group Dependency view.

See “Viewing service group dependencies” on page 113.

Administering resources in a cluster

The Cluster Management Console enables you to add and configure resources according to the requirements of the service groups that the resource supports. Use the following views to administer a resource:

- Cluster:Resources
  This is a cluster-level view of all resources in the cluster.
- Resource:Summary
  This is the first view at the resource level in the VCS Management Console.
- Group:Summary
  This is the first view at the service group level in the VCS Management Console.
- Group:Resources
  This is a service-group-level view of all resources in a service group and any dependencies that exist among them.

To navigate to the Cluster:Resources view

1. On the navigation bar, click **[*Cluster*]**.
2. On the secondary tab bar, click **Summary** and then click **Resources**.

In the **Cluster:Resources** view, the **Resources Listing** table lists all resources in the selected cluster. The task panels contain tasks that you can perform in a one-to-many fashion on the listed resources.

To navigate to the Resource:Summary view

1. On the navigation bar, click **[*Cluster*]**.
2. On the secondary tab bar, click **Summary** and then click **Resources**.
3. In the Cluster:Resources view, in the Resources Listing table, click a linked resource name.
   In the **Resource:Summary** view, you can choose a task from the task pane or select another resource-level view using the tabs on the secondary tab bar.
   Each view contains information and tasks for administering a single resource.

To navigate to the Group:Summary view

1. On the navigation bar, click **[*Cluster*]**.
2. On the secondary tab bar, click **Summary** and then click **Groups**.
3. In the **Cluster:Groups** view, in the **Service Groups Listing** table, click a linked service group name.
   In the **Group:Summary** view, the **Configuration** task panel contains several resource-related tasks.

To navigate to the Group:Resources view

1. On the navigation bar, click **[*Cluster*]**.
2. On the secondary tab bar, click **Summary** and then click **Groups**.
3. In the Cluster:Groups view, in the Service Groups Listing table, click a linked service group name.
4. In the Group:Summary view, on the secondary tab bar, click **Resources**.
   The Group:Resources view contains a graphical representation of the resources in the service group and any dependencies that exist among them.

Adding a resource to a service group

Add and configure a required resource as a member of a service group. You must have the role of cluster administrator or service group administrator to add a resource to a service group.

To add a resource to a service group

1 In the Group:Summary view, in the Configuration task panel, click **Add/Modify Resource**.

2 In the Resource Creation dialog box, specify the following details about the resource and the service group and then click **OK**:

- The resource name

  Enter the name of the resource that you want to add to the service group into the **Resource Name** text box. Names must consist of only one string and must not contain special characters like . (period) or _ (underscore).

- The resource type

  Select the type of the new resource from the **Resource Type** drop-down menu.

- The resource startup state

  If you want the resource to start in an enabled state after the service group comes online, check **Enable Resource**.

- The resource list for the service group

  Click **Add Resource** to add the resource to the member resource list for the service group. The resource list is used to populate the **Resource List** table.

  In the Resource List table, you can edit resource attributes or delete a resource using the buttons in the Edit and Delete columns.

Overriding resource attributes

Change the default value of one or more resource attributes. You must have the role of cluster administrator or group administrator to override resource attribute values.

To override resource attributes

1 In the Cluster:Resources view, in the Resources Listing table, click a linked resource name.

2 In the Resource:Summary view, in the Configuration task panel, click **Override Attributes**.

3 In the Override Attributes dialog box, select the attribute you want to override from the drop-down menu and then click **OK**.

4 In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:

- The attribute value
  In the Values field, either click a value to change the it, or click the **+** (plus) button to add a value. Click a value and then click the **-** (minus) button to remove a value.
  You can enter more than one value for some attributes. Use the up- and down-arrow buttons to scroll among multiple values.
- The systems on which to change the attribute
  Specify a setting for the Apply value to option. **Apply to all nodes** applies the changes to all systems listed the System List, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the System List.

To remove resource overrides

1. In the Cluster:Resources view, in the Resources Listing table, click a linked resource name.
2. In the Resource:Summary view, in the Configuration task panel, click **Remove Overrides**.
3. Click **OK** to confirm that you want to remove previously configured override values.

Enabling and disabling resources

Prepare a resource to be brought online or prevent a resource from being brought online. A resource must be enabled before it can be brought online. A disabled resource cannot be brought online. You must have the role of cluster operator or service group operator to enable or disable resources.

To enable or disable selected resources

1. In the Cluster:Resources view, in the Resources Listing table, click a linked resource name.
2. In the Operations task panel, click **Enable** or **Disable**.
3. In either the Enable Resource or Disable Resource dialog box, click **OK**.

Bringing a resource online

Manually bring a resource online. You must have the role of cluster operator or group operator to bring a resource online.

1. In the Cluster:Resources view, in the Resources Listing table, click the linked name of the resource you want to bring online.
2. In the Resource:Summary view, in the Operations task panel, click **Online**.

3 In the Online Resource dialog box, select the system on which to bring the resource online from the drop-down menu.

4 Click **OK**.

Taking a resource offline

Manually take a resource offline. You must have the role of cluster operator or group operator to take a resource online.

1 In the Cluster:Resources view, in the Resources Listing table, click the linked name of the resource you want to bring online.

2 In the Resource:Summary view, in the Operations task panel, click **Online**.

3 In the Online Resource dialog box, select the system on which to bring the resource online from the drop-down menu.

4 Click **OK**.

Taking a resource offline and propagating the command

Take a resource and all of its dependents offline. Use this task to propagate the offline state of a parent resource to its child resources. This link is disabled if any of the following conditions exist:

- You do not have the role of either cluster administrator or cluster operator
- The resource does not depend on any other resource
- An online resource depends on this resource
- The resource is not online

To take a parent resource and all of its child resources offline

1 In the Cluster:Resources view, in the Resources Listing table, click the linked name of the resource you want to take offline.

2 In the Resource:Summary view, in the Operations task panel, click **Offline Propagate**.

3 In the Offline Propagate for Resource dialog box, select the system on which to bring the resource online from the drop-down menu.

4 Click **OK**.

Faulting a resource

Manually simulate a fault on a resource to test failover behavior. This operation is available only when using the cluster simulator.

To fault a single service group

1. In the Cluster:Resources view, in the Resources Listing table, click the linked name of the resource that you want to fault.
2. In the Resource:Summary view, in the Simulation task panel, click **Fault Resource**.
3. In the Fault Resource dialog box, on the drop-down menu, select the systems on which you want to fault the resource.
   This drop-down menu is populated using the system list of the resource. The system list of a resource contains a list of all systems on which the resource is configured.
4. Click **OK**.

To fault one or more resources

1. In the Cluster:Resources view, in the Resources Listing table, check the check box preceding the line item for each resource that you want to fault. Select *all* resources by checking the check box at the top.
2. In the Simulation task panel, click **Fault Resources**.
   In the Fault Resource(s) dialog box, use the appropriate drop-down menu to select the systems on which you want to fault each resource. Each drop-down menu is populated using the system list of each resource. The system list of a resource contains a list of all systems on which the resource is configured.
3. Click **OK**.

Clearing a faulted resource

Clear a resource to remove a fault and make the resource available to go online. You must have the role of cluster operator or service group operator to clear a faulted resource. A resource fault can result from occurences such as a power failure or a faulty configuration.

To clear a single faulted resource

1. In the Cluster:Resources view, in the Resources Listing table, click the linked name of the faulted resource that you want to clear.
2. In the Resource:Summary view, in the Operations task panel, click **Clear Fault**.

To clear one or more faulted resources

1. In the Cluster:Resources view, in the Resources Listing table, check the check box preceding the line item for each resource that you want to clear. Select all resources by checking the check box at the top.
2. In the Operations task panel, click **Clear Fault**.

Probing a resource

Probe a resource to confirm that it is properly configured and ready to bring online.

To probe a resource

1. In the Cluster:Resources view, in the Resources Listing table, click the linked name of the resource that you want to probe.
2. In the Resource:Summary view, in the Operations task panel, click **Probe**.
3. In the Probe Resource dialog box, select the system on which you want to probe the specified resource from the drop-down menu and then click **OK**.
   The dialog box displays a message indicating success or failure. Success indicates that the resource is properly configured and, if not already online, is ready to be brought online.

Deleting a resource from a service group

Delete a resource from membership in a service group. You must have the role of cluster administrator or service group administrator to delete a resource from a service group.

To delete a resource from a service group

1. In either the Cluster:Resources or Group:Summary view, in the Resource Listing table, click the linked name of the resource that you want to delete.
2. In the Resource:Summary view, in the Configuration task panel, click **Delete Resource**.
3. In the Delete Resource dialog box, click **OK** to confirm that you want to delete the resource.

Linking resources

Link resources to create a dependency. You must have the role of cluster administrator or service group administrator to link resources.

To link two resources

1 In the Resource:Summary view, in the Configuration task panel, click **Link/Unlink Resource**.

2 In the Link/Unlink Resource dialog box, specify the following details for the resource link and then click **OK**:

- The parent resource
  Select the parent resource from the drop-down menu. This resource must be brought online before the child resource. The **Parent Resource** menu is populated using the resource list of the service group in which the resources are configured.
- The child resource
  Select the child resource from the scroll box. This resource is brought online after the parent resource. The **Child Resource** scroll box is populated using the resource list of the service group in which the resources are configured. This list is filtered depending upon the selection you make in **Parent Resource**. Click **Add Link** after you select a child resource to add the link to the **Currently Linked** table.

Viewing resource dependencies

View a graphical representation of the dependencies that exist for a resource.

To view service group dependencies and get status information

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Summary** and then click **Resources**.

3 In the Resources Listing table, click the linked name of a resource.

4 In the Resource:Summary view, on the secondary tab bar, click **Dependency**.
This view presents a graphical representation of the selected resource and all currently configured resources that have a dependency relationship with the resource.

5 Click a resource icon to get information about the resource that it represents.

Unlinking resources

Unlink two resources to remove a dependency. You must have the role of cluster administrator or service group administrator to unlink resources.

To unlink two resources

1. In the Resource:Summary view, in the Configuration task panel, click **Link/Unlink Resource**.
2. In the Link/Unlink Resource dialog box, in the Currently Linked table, click the **X** button on the line item that corresponds to the link that you want to remove.
3. Click **OK**.

Modifying a resource attribute

Edit the attributes of a resource. You must have the role of cluster administrator or service group administrator to edit resource attributes.

To edit a resource attribute using the attribute table

1. In the Resource:Summary view, on the secondary tab bar, click **Attributes**.
2. In the All attributes for resource table, locate the line item for the attribute you want to edit.
3. In the **Edit** column, click the **...** (edit) button.
4. In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:
   - The attribute value
     In the Values field, either click a value to change the it, or click the **+** (plus) button to add a value. Click a value and then click the **-** (minus) button to remove a value.
     You can enter more than one value for some attributes. Use the up- and down-arrow buttons to scroll among multiple values.
   - The systems on which to change the attribute
     Specify a setting for the **Apply value to** option. **Apply to all nodes** applies the changes to all systems listed the **System List**, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the **System List**.

Invoking a resource action

Run a predefined action script on a resource. Examples of predefined resource actions include splitting and joining disk groups.

To invoke a resource action

1. In the Cluster:Resources view, in the Resources Listing table, click the linked name of the resource on which you want to invoke an action.

2 In the Resource:Summary view, in the Operations task panel, click **Invoke Action**.

3 In the Invoke Action dialog box, select the script that you want to run and the system on which you want to run it.

4 Specify arguments for the action
Click the **+** (plus) button to add arguments. Click an argument name and then click the **–** (minus) button to remove it.

5 Click **OK**.

Refreshing the ResourceInfo attribute

Refresh the ResourceInfo attribute to view the latest values for that attribute.

To refresh the ResourceInfo attribute

1 In the Cluster:Resources view, in the Resources Listing table, click the linked name of a resource with the ResourceInfo attribute.

2 In the Resource:Summary view, in the Operations task panel, click **Refresh Resource Info**.

3 In the Refresh Resource Info dialog box, select the system on which you want to refresh the ResourceInfo attribute.

4 Click **OK**.

Clearing the ResourceInfo attribute

Clear the ResourceInfo attribute to reset all the parameters in this attribute.

To clear the parameters of the ResourceInfo attribute

1 In the Cluster:Resources view, in the Resources Listing table, click the linked name of a resource with the ResourceInfo attribute.

2 In the Resource:Summary view, in the Operations task panel, click **Clear Resource Info**.

3 In the Clear Resource Info dialog box, select the system on which you want to clear the ResourceInfo attribute.

Administering resource types in a cluster

The Cluster Management Console enables you to add and configure resource types according to the requirements of the resources that they represent. Use the following views to administer a resource type:

- Cluster:Resource Types
  This is a cluster-level view of all resource types in the cluster.
- Resource Type:Attributes
  This is the first and only view at the resource type level in the VCS Management Console.

To navigate to the Cluster:Resource Types view

1. On the navigation bar, click **[*Cluster*]**.
2. On the secondary tab bar, click **Summary** and then click **Resource Types**.
   In the **Cluster:Resource Types** view, the **Resource Types Listing** table lists all resource types in the selected cluster. The task panels contain tasks that you can perform in a one-to-many fashion on the listed resource types.

To navigate to the Resource Type:Attributes view

1. On the navigation bar, click **[*Cluster*]**.
2. On the secondary tab bar, click **Summary** and then click **Resource Types**.
3. In the Cluster:Resource Types view, in the Resource Types Listing table, click a linked resource type name.
   In the Resource Type:Attributes view, you can choose a task from the task pane or click the button in the **Edit** column of the table to edit the attribute value.

The only management task available is to edit the resource type attributes. You must have the role of cluster administrator or service group administrator to edit resource type attributes.

To edit resource type attributes using the attribute table

1. In the Cluster:Resource Types view, in the Resource Types Listing table, click a linked resource name.
2. In the Resource Type:Attributes view, in the All attributes for resource type table, locate the line item for the attribute that you want to edit. Not all attributes can be edited.
3. In the Edit column, click the **...** (edit) button.
4. In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:

- The attribute value
  In the Values field, specify an appropriate string value for the attribute.
- The systems on which to change the attribute
  Specify a setting for the **Apply value to** option. **Apply to all nodes** applies the changes to all systems listed the System List, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the System List.

Administering systems in a cluster

The VCS Management Console enables you to configure and manage the systems that host the cluster service groups and resources. Use the **Cluster:Systems** view and the **System:Summary** view to administer systems.

To navigate to the Cluster:Systems view

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Systems**.
In the Cluster:Systems view, the Systems Listing table lists all systems in the selected cluster. The task panels contain several tasks that you can perform in a one-to-many fashion on the listed systems.

To navigate to the System:Summary view

1 On the navigation bar, click **[*Cluster*]**.

2 On the secondary tab bar, click **Systems**.

3 In the Cluster:Systems view, in the Systems Listing table, click a linked system name.
In the Systems:Summary view, you can choose a task from the task pane or select another system-level view using the tabs on the secondary tab bar.
Each view contains information and tasks for administering a single system.

Freezing a System

Freeze a system to prevent the service groups that it hosts from failing over to another system. Use this procedure before performing a system upgrade.

To freeze a system

1 In the [Cluster]:Systems view, in the Systems Listing table, click the linked name of the system that you want to freeze.

2 In the System:Summary view, in the Operations task panel, click **Freeze**.

3 In the Freeze System dialog box, select one or both of the following options and then click **OK**:
   - Select **Persistent** if you want the system to retain a frozen state when the cluster is rebooted.
   - Select **Evacuate** if you want the system's active service groups to fail over to another system in the cluster before the freezing operation takes place.

You can freeze one or more systems at a time directly from the Cluster:Systems view by selecting the check boxes for systems in the Systems Listing table and then clicking **Freeze** in the **Operations** task panel.

Unfreezing a system

Unfreeze a frozen system to perform online or offline operations on the system.

To unfreeze a system

1 In the [Cluster]:Systems view, in the Systems Listing table, click the linked name of the system that you want to unfreeze.

2 In the System:Summary view, in the Operations task panel, click **Unfreeze**.

3 In the Unfreeze System dialog box, click **OK** to confirm that you want to unfreeze the system.

You can unfreeze one or more systems at a time directly from the Cluster:Systems view by selecting the check boxes for systems in the Systems Listing table and then clicking **Unfreeze** in the **Operations** task panel.

Modifying system attributes

Edit the attributes of a system. You must have the role of cluster administrator or service group administrator to edit system attributes.

Edit the attributes of a cluster. You must have the role of cluster administrator to edit cluster attributes.

To edit a cluster attribute using the attribute table

1 In the System:Summary view, on the secondary tab bar, click **Attributes**.

2 In the All attributes for system table, locate the line item for the attribute you want to edit. Not all attributes can be edited.

3 In the Edit column, click the **...** (edit) button.

4 In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:

- The attribute value
  In the Values field, either click a value to change the it, or click the **+** (plus) button to add a value. Click a value and then click the **-** (minus) button to remove a value.
  You can enter more than one value for some attributes. Use the up- and down-arrow buttons to scroll among multiple values.
- The systems on which to change the attribute
  Specify a setting for the **Apply value to** option. **Apply to all nodes** applies the changes to all systems listed the System List, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the System List.

To edit a system attribute using the task pane

1. In the Cluster:Summary view, in the Clusters Listing table, check the check box preceding the line item for each system with the attribute that you want to edit. Select all systems by checking the check box at the top.
2. In the Configuration task panel, click **Edit Attribute**.
3. Specify the attribute details.
   See "Modifying system attributes" on page 130.

Administering attributes in a cluster

The VCS Management Console enables you to edit the attributes of cluster objects (clusters, service groups, systems, resources, and resource types). Ensure that the cluster configuration is open (in read/write mode) before editing attributes. Attributes for an object are available in any object summary view, such as the Cluster:Summary view or the Resource:Summary view. After navigating to an object summary view, click the Attributes tab on the secondary tab bar.

Changes to certain attributes, such as a webip attribute, may involve taking the service group offline, modifying the configuration file, and bringing the group back online.

To edit an attribute

1. Navigate to an *Object*:Summary view for the object with the attributes you want to edit.
2. In the *Object*:Summary view, on the secondary tab bar, click **Attributes**.
3. In the All attributes table, click the **...** (edit) button for the attribute you want to edit.

4 In the Edit Attribute dialog box, specify the following attribute details and then click **OK**:

- The attribute value

  In the Values field, either click a value to change the it, or click the **+** (plus) button to add a value. Click a value and then click the **-** (minus) button to remove a value. (Some Edit Attribute dialog boxes have just a text field for the value).

  You can enter more than one value for some attributes. Use the up- and down-arrow buttons to scroll among multiple values.

- The systems on which to change the attribute

  Specify a setting for the **Apply value to** option. **Apply to all nodes** applies the changes to all systems listed the System List, which lists all systems on which the attribute is configured. **Apply to selected nodes** enables you to select nodes from the System List.

Conducting a search

The search feature uses text-based comparison. You can perform a search using either the search bar or the Search tab.

The search bar is located immediately below the main tab bar, and contains a text entry box and a Search button. (You may need to scroll right to see these controls on some systems.) The search bar is available for all views except those under the Search tab.

The Search tab is located on the main tab bar. Views under the Search tab contain a text entry box and a Search button and offer some filtering options that the search bar does not.

The search feature displays results in one or more of the search views, which you navigate to using secondary tabs under Search. These tabs are labeled Clusters, Groups, Systems, Resources, Resource Types, and Logs. Each tab contains a table that lists the search results, or hits, in the corresponding object category. By default, the console takes you to the search views after each search. The destination view is whichever view was most recently visited.

Searches are limited to text strings. Separate two or more strings using a space character. The search feature does not currently support:

- Direct textual entry of Boolean operators
- Direct textual entry of wildcards
- Query expressions

To perform a search using the search bar

1. In the text entry box, type one or more text strings that you want to find.
   Use a space character to separate text strings.
2. Click **Search**.
3. Visit each search view to see the hits in each object category.
   Your particular destination view depends on which search view was most recently visited. Click the other secondary tabs (**Clusters**, **Groups**, and so on) to visit the other search views.

To perform a search using the Search tab

1. On the main tab bar, click **Search**.
2. On the secondary tab bar, click the object category (**Clusters**, **Groups**, and so on) in which you want to begin your search.
3. In the text entry box, type one or more text strings that you want to find.
   Use a space character to separate text strings.

4 Check one or more of the following options and then click **Search**:

- **Exact Match**

  Requires each search result to match at least one whole search string. This option excludes occurrences in which a search string is part of a longer string.

- **Match All Terms**

  Requires each search result to match all search terms (strings).

  The behavior of this option depends on the setting of the **Exact Match** option according to the following example:

  Consider object1 with attribute path = /tmp/f3 and object2 with attribute path = /tmp/f4:

Search String	Exact Match Option Setting	Match All Terms Option Setting	Search results
path f3	checked	checked	none
path f3	checked	cleared	none
path f3	cleared	checked	object1
path f3	cleared	cleared	object1, object2

- **Highlight**

  Highlights the string occurrences in the search result tables.

Combine **Exact Match** and **Match All Terms** to minimize the number of candidate occurrences.

5 Visit each search view to see the hits in each object category.

Your first results are displayed in the same view in which you started the search. Click the other secondary tabs (**Clusters**, **Groups**, and so on) to see any hits the other search views.

Note: After you perform a search, changing the Exact Match, Match All Terms, and Highlight option settings automatically updates the results.

Chapter 7

Administering the cluster from Cluster Manager (Java console)

- About the Cluster Manager (Java Console)
- Getting started
- Reviewing components of the Java Console
- About Cluster Monitor
- About Cluster Explorer
- Accessing additional features of the Java Console
- Administering Cluster Monitor
- Administering user profiles
- Administering service groups
- Administering resources
- Administering systems
- Administering clusters
- Executing commands
- Editing attributes
- Querying the cluster configuration
- Setting up VCS event notification using the Notifier wizard
- Administering logs
- Administering VCS Simulator

About the Cluster Manager (Java Console)

Cluster Manager (Java Console) offers complete administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types. Many of the operations supported by the Java Console are also supported by the command line interface and Cluster Management Console.

The console enables or disables features depending on whether the features are supported in the cluster that the console is connected to. For example, the Cluster Shell icon is grayed out when you connect to recent versions of VCS. But the icon is enabled when you connect to earlier versions of a VCS cluster.

Getting started

- Make sure you have the current version of Cluster Manager (Java Console) installed. If you have a previous version installed, upgrade to the latest version. Cluster Manager (Java Console) is compatible with earlier versions of VCS.
- Cluster Manager (Java Console) is supported on:
  - Windows 2000, Windows XP, and Windows 2003

  Note: Make sure you are using the operating system version that supports JRE 1.5.

- Verify the configuration has a user account. A user account is established during VCS installation that provides immediate access to Cluster Manager. If a user account does not exist, you must create one.
  See "Adding a user" on page 173.
- Start Cluster Manager.
  See "Starting Cluster Manager (Java console)" on page 137.
- Add a cluster panel.
  See "Configuring a new cluster panel" on page 168.
- Log on to a cluster.
  See "Logging on to and off of a cluster" on page 170.
- Make sure you have adequate privileges to perform cluster operations.
  See "About the VCS user privilege model" on page 63.

Starting Cluster Manager (Java console)

To start the Java Console on Windows systems

Double-click the Veritas Cluster Manager (Java Console) icon on the desktop.

Reviewing components of the Java Console

Cluster Manager (Java Console) offers two windows, Cluster Monitor and Cluster Explorer, from which most tasks are performed. Use Cluster Manager to manage, configure, and administer the cluster while VCS is running (online).

The Java Console also enables you to use VCS Simulator. Use this tool to simulate operations and generate new configuration files (main.cf and types.cf) while VCS is offline. VCS Simulator enables you to design configurations that imitate real-life scenarios without test clusters or changes to existing configurations.

See “Administering VCS Simulator” on page 234.

Icons in the Java Console

The Java Console uses the following icons to communicate information about cluster objects and their states.

See “Cluster and system states” on page 693.

Table 7-4 Icons in Cluster Manager (Java Console)

Icon	Description
	Cluster
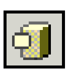	System
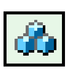	Service Group
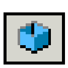	Resource Type
	Resource
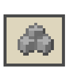	OFFLINE
	Faulted (in UP BUT NOT IN CLUSTER MEMBERSHIP state)

Table 7-4 Icons in Cluster Manager (Java Console)

Icon	Description
	Faulted (in EXITED state)
	PARTIAL
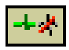	Link Heartbeats (in UP and DOWN states)
	UP AND IN JEOPARDY
	FROZEN
	AUTODISABLED
	UNKNOWN
	ADMIN_WAIT
	Global Service Group (requires the VCS Global Cluster Option)
	Remote Cluster in RUNNING state (requires the VCS Global Cluster Option)
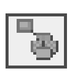	Remote Cluster in EXITING, EXITED, INIT, INQUIRY, LOST_CONN, LOST_HB, TRANSITIONING, or UNKNOWN state.

About Cluster Monitor

After starting Cluster Manager, the first window that appears is Cluster Monitor. This window includes one or more panels displaying general information about actual or simulated clusters. Use Cluster Monitor to log on to and off of a cluster, view summary information on various VCS objects, customize the display, use VCS Simulator, and exit Cluster Manager.

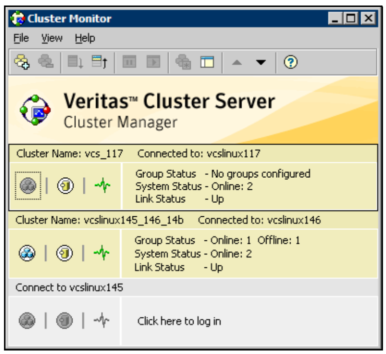

Cluster monitor toolbar

The Cluster Monitor toolbar contains the following buttons.

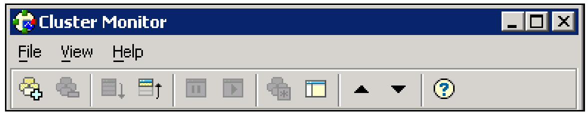

From left to right:


New Cluster. Adds a new cluster panel to Cluster Monitor.

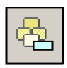
Delete Cluster. Removes a cluster panel from Cluster Monitor.

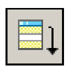
Expand. Expands the Cluster Monitor view.


Collapse. Collapses the Cluster Monitor view.

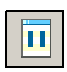
Stop. Pauses cluster panel scrolling.

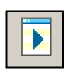
Start. Resumes scrolling.

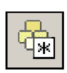
Login. Log on to the cluster shown in the cluster panel.

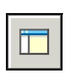
Show Explorer. Launches an additional window of Cluster Explorer after logging on to that cluster.

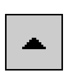
Move Cluster Panel Up. Moves the selected cluster panel up.

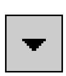
Move Cluster Panel Down. Moves the selected cluster panel down.


Help. Access online help.

Cluster monitor panels

To administer a cluster, add a cluster panel or reconfigure an existing cluster panel in Cluster Monitor. Each panel summarizes the status of the connection and components of a cluster.

Monitoring the cluster connection with Cluster Monitor

The right pane of a panel in Cluster Monitor displays the status of the connection to a cluster. An inactive panel will appear grey until the user logs on and connects to the cluster. To alter the connection to a cluster, right-click a panel to access a menu.

- The menu on an active panel enables you to log off a cluster.
- The menu on an inactive panel enables you to log on to a cluster, configure the cluster, and delete the cluster from Cluster Monitor.

Menus are enabled when the Cluster Monitor display appears in the default expanded view. If you activate a menu on a collapsed scrolling view of Cluster Monitor, the scrolling stops while accessing the menu.

If the system to which the console is connected goes down, a message notifies you that the connection to the cluster is lost. Cluster Monitor tries to connect to another system in the cluster according to the number of Failover retries set in the Connectivity Configuration dialog box. The panels flash until Cluster Monitor is successfully connected to a different system. If the failover is unsuccessful, a message notifies you of the failure and the panels turn grey.

Monitoring VCS objects with Cluster Monitor

Cluster Monitor summarizes the state of various objects in a cluster and provides access to in-depth information about these objects in Cluster Explorer. The right pane of a Cluster Monitor panel displays the connection status (online, offline, up, or down) of service groups, systems, and heartbeats. The left pane of a Cluster Monitor panel displays three icons representing service groups, systems, and heartbeats. The colors of the icons indicate the state of the cluster; for example:

- A flashing red slash indicates Cluster Manager failed to connect to the cluster and will attempt to connect to another system in the cluster.
- A flashing yellow slash indicates Cluster Manager is experiencing problems with the connection to the cluster.

Pointing to an icon accesses the icon’s ScreenTip, which provides additional information on the specific VCS object.

To review detailed information about VCS objects in Cluster Explorer, Logs, and Command Center, right-click a panel to access a menu. Menus are enabled when the Cluster Monitor display appears in the default expanded view. If you activate a menu on a collapsed scrolling view of Cluster Monitor, the scrolling stops while accessing the menu.

Expanding and collapsing the Cluster Monitor display

Cluster Monitor supports two views: expanded (default) and collapsed. The expanded view shows all cluster panels. The collapsed view shows one cluster panel at a time as the panels scroll upward.

Operations enabled for the expanded view of cluster panels, such as viewing menus, are also enabled on the collapsed view after the panels stop scrolling.

To collapse the Cluster Monitor view

On the **View** menu, click **Collapse**.

or

Click **Collapse** on the Cluster Monitor toolbar.

To expand the Cluster Monitor view

On the **View** menu, click **Expand**.

or

Click **Expand** on the Cluster Monitor toolbar.

To pause a scrolling cluster panel

Click the cluster panel.

or

Click **Stop** on the Cluster Monitor toolbar.

Customizing the Cluster Manager display

Customize the Cluster Manager to display objects according to your preference.

To customize the Cluster Manager display

1 From Cluster Monitor, click **Preferences** on the **File** menu. If you are using a Windows system, proceed to step 2. Otherwise, proceed to step 3.

2 In the **Look & Feel** tab (for Windows systems):

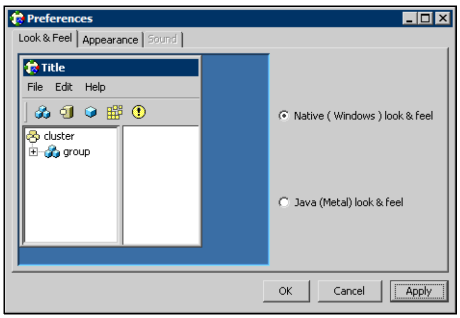

- Click **Native (Windows** or **Motif) look & feel** or **Java (Metal) look & feel**.
- Click **Apply**.

3 In the **Appearance** tab:

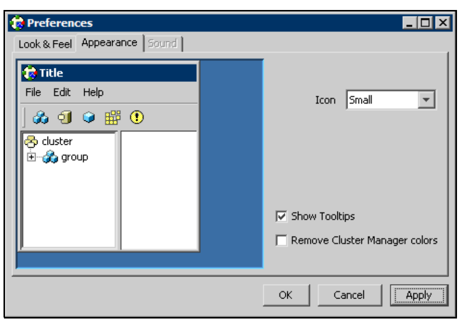

- Click the color (applies to Java (Metal) look & feel).
- Click an icon size.
- Select the **Show Tooltips** check box to enable ToolTips.
- Select the **Remove Cluster Manager colors** check box to alter the standard color scheme.
- Click **Apply**.

4 In the **Sound** tab:

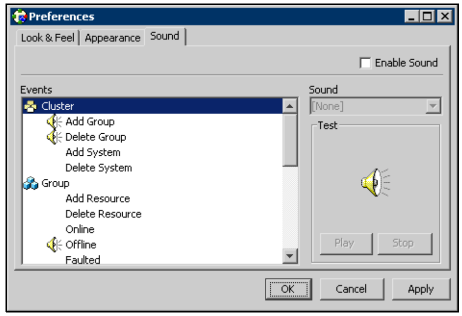

Note: This tab requires a properly configured sound card.

- Select the **Enable Sound** check box to associate sound with specific events.
- Click an event from the **Events** configuration tree.
- Click a sound from the **Sounds** list box.
- To test the selected sound, click **Play**.
- Click **Apply**.
- Repeat these steps to enable sound for other events.

5 After you have made your final selection, click **OK**.

About Cluster Explorer

Cluster Explorer is the main window for cluster administration. From this window, you can view the status of VCS objects and perform various operations.

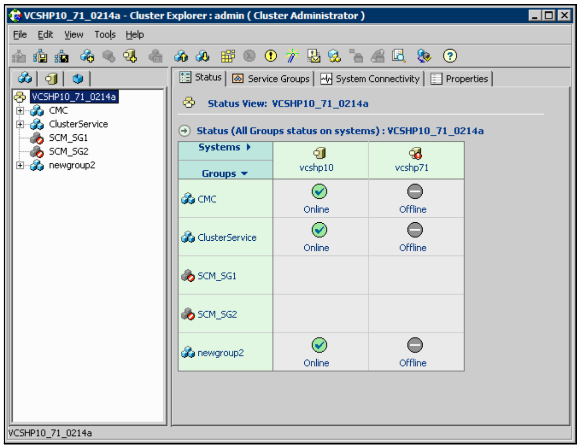

The display is divided into three panes. The top pane includes a toolbar that enables you to quickly perform frequently used operations. The left pane contains a configuration tree with three tabs: Service Groups, Systems, and Resource Types. The right pane contains a panel that displays various views relevant to the object selected in the configuration tree.

To access Cluster Explorer

1. Log on to the cluster.
2. Click anywhere in the active Cluster Monitor panel.
   or
   Right-click the selected Cluster Monitor panel and click Explorer View from the menu.

Cluster Explorer toolbar

The Cluster Explorer toolbar contains 18 buttons. Available operations are described below. Note: Some buttons may be disabled depending on the type of cluster (local or global) and the privileges with which you logged on to the cluster.

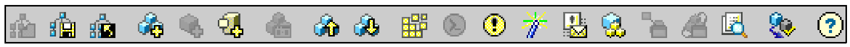

From left to right:

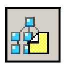
Open Configuration. Modifies a read-only configuration to a read-write file. This enables you to modify the configuration.

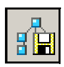
Save Configuration. Writes the configuration to disk.

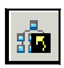
Save and Close Configuration. Writes the configuration to disk as a read-only file.


Add Service Group. Displays the Add Service Group dialog box.


Add Resource. Displays the Add Resource dialog box.

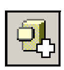
Add System. Displays the Add System dialog box.

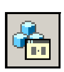
Manage systems for a Service Group. Displays the System Manager dialog box.


Online Service Group. Displays the Online Service Group dialog box.

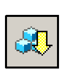
Offline Service Group. Displays the Offline Service Group dialog box.

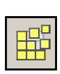
Show Command Center. Enables you to perform many of the same VCS operations available from the command line.

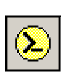
Show Shell Command Window. Enables you to launch a non-interactive shell command on cluster systems, and to view the results on a per-system basis.

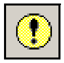
Show the Logs. Displays alerts and messages received from the VCS engine, VCS agents, and commands issued from the console.

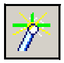
Launch Configuration Wizard. Enables you to create VCS service groups.


Launch Notifier Resource Configuration Wizard. Enables you to set up VCS event notification.


Remote Group Resource Configuration Wizard. Enables you to configure resources to monitor a service group in a remote cluster.


Add/Delete Remote Clusters. Enables you to add and remove global clusters.


Configure Global Groups. Enables you to convert a local service group to a global group, and vice versa.


Query. Enables you to search the cluster configuration according to filter criteria.


Show Cluster Explorer Help. Enables you to access online help.

Cluster Explorer configuration tree

The Cluster Explorer configuration tree is a tabbed display of VCS objects.

- The **Service Groups** tab lists the service groups in the cluster. Expand each service group to view the group's resource types and resources.
- The **Systems** tab lists the systems in the cluster.
- The **Types** tab lists the resource types in the cluster

Cluster Explorer view panel

The right pane of the Cluster Explorer includes a view panel that provides detailed information about the object selected in the configuration tree. The information is presented in tabular or graphical format. Use the tabs in the view panel to access a particular view. The console enables you to "tear off" each view to appear in a separate window.

- Click any object in the configuration tree to access the Status View and Properties View.
- Click a cluster in the configuration tree to access the Service Group view, System Connectivity view, and Remote Cluster Status View (for global clusters only).
- Click a service group in the configuration tree to access the Resource view.

To create a tear-off view

On the **View** menu, click **Tear Off**, and click the appropriate view from the menu.

or

Right-click the object in the configuration tree, click **View**, and click the appropriate view from the menu.

Status view

The Status View summarizes the state of the object selected in the configuration tree. Use this view to monitor the overall status of a cluster, system, service group, resource type, and resource.

For example, if a service group is selected in the configuration tree, the Status View displays the state of the service group and its resources on member systems. It also displays the last five critical or error logs. Point to an icon in the status table to open a ScreenTip about the relevant VCS object.

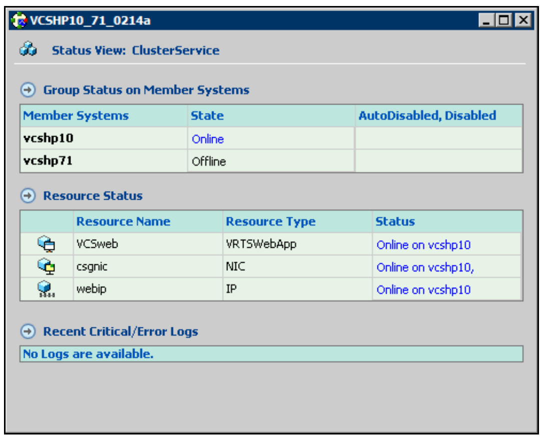

For global clusters, this view displays the state of the remote clusters. For global groups, this view shows the status of the groups on both local and remote clusters.

To access the Status view

1. From Cluster Explorer, click an object in the configuration tree.
2. In the view panel, click the **Status** tab.

Properties view

The Properties View displays the attributes of VCS objects. These attributes describe the scope and parameters of a cluster and its components.

To view information on an attribute, click the attribute name or the icon in the **Help** column of the table.

See “VCS attributes” on page 701”

By default, this view displays key attributes of the object selected in the configuration tree. The Properties View for a resource displays key attributes of the resource and attributes specific to the resource types. It also displays attributes whose values have been overridden.

See “Overriding resource type static attributes” on page 209.

To view all attributes associated with the selected VCS object, click **Show all attributes**.

To access the properties view

1. From Cluster Explorer, click a VCS object in the configuration tree.
2. In the view panel, click the **Properties** tab.

Service Group view

The Service Group view displays the service groups and their dependencies in a cluster. Use the graph and ScreenTips in this view to monitor, create, and disconnect dependencies. To view the ScreenTips, point to a group icon for information on the type and state of the group on the cluster systems, and the type of dependency between the service groups.

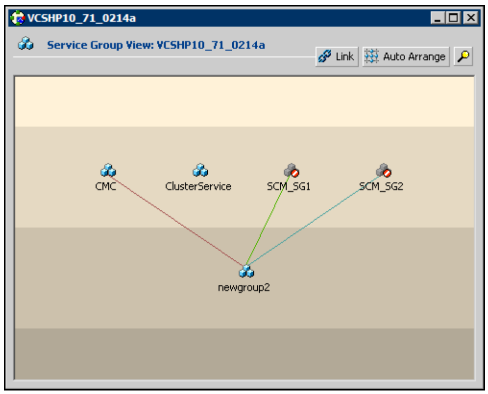

The line between two service groups represents a dependency, or parent-child relationship. In VCS, parent service groups depend on child service groups. A service group can function as a parent and a child.

See "About service group dependencies" on page 492.

The color of the link between service groups indicates different types of dependencies.

- A blue link indicates a soft dependency.
- A red link indicates a firm dependency.
- A green link indicates a hard dependency typically used with VVR in disaster recovery configurations.

To access the Service Group view

1. From Cluster Explorer, click a cluster in the configuration tree.
2. In the view panel, click the **Service Groups** tab.

Resource view

The Resource view displays the resources in a service group. Use the graph and ScreenTips in this view to monitor the dependencies between resources and the status of the service group on all or individual systems in a cluster.

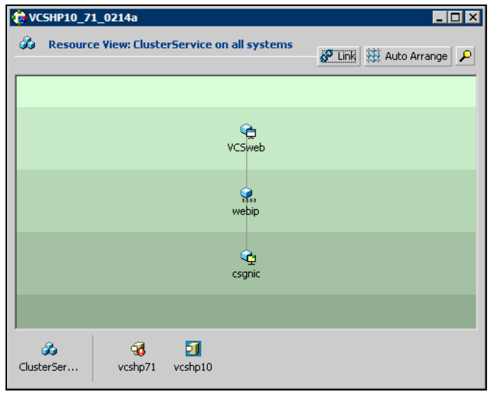

In the graph, the line between two resources represents a dependency, or parent-child relationship. Resource dependencies specify the order in which resources are brought online and taken offline. During a failover process, the resources closest to the top of the graph must be taken offline before the resources linked to them are taken offline. Similarly, the resources that appear closest to the bottom of the graph must be brought online before the resources linked to them can come online.

- A resource that depends on other resources is a parent resource. The graph links a parent resource icon to a child resource icon below it. Root resources (resources without parents) are displayed in the top row.
- A resource on which the other resources depend is a child resource. The graph links a child resource icon to a parent resource icon above it.
- A resource can function as a parent and a child.

Point to a resource icon to display ScreenTips about the type, state, and key attributes of the resource. The state of the resource reflects the state on a specified system (local).

In the bottom pane of the Resource view, point to the system and service group icons to display ScreenTips about the service group status on all or individual systems in a cluster. Click a system icon to view the resource graph of the service group on the system. Click the service group icon to view the resource graph on all systems in the cluster.

To access the Resource view

1. From Cluster Explorer, click the service groups tab in the configuration tree.
2. Click a service group in the configuration tree.
3. In the view panel, click the **Resources** tab.

Moving and linking icons in Service Group and Resource views

The Link and Auto Arrange buttons are available in the top right corner of the Service Group or Resource view:

Click **Link** to set or disable the link mode for the Service Group and Resource views.

Note: There are alternative ways to set up dependency links without using the Link button.

The link mode enables you to create a dependency link by clicking on the parent icon, dragging the yellow line to the icon that will serve as the child, and then clicking the child icon. Use the Esc key to delete the yellow dependency line connecting the parent and child during the process of linking the two icons.

If the Link mode is *not* activated, click and drag an icon along a horizontal plane to move the icon. Click **Auto Arrange** to reset the appearance of the graph. The view resets the arrangement of icons after the addition or deletion of a resource, service group, or dependency link. Changes in the Resource and Service Group views will be maintained after the user logs off and logs on to the Java Console at a later time.

Zooming in on Service Group and Resource views

The Resource view and Service Group view include a navigator tool to zoom in or out of their graphs. Click the magnifying glass icon in the top right corner to open the zoom panel.

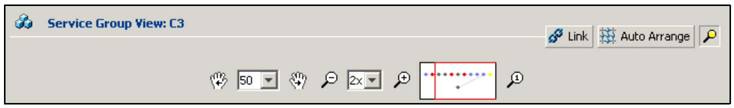

- To move the view to the left or right, click a distance (in pixels) from the drop-down list box between the hand icons. Click the <- or -> hand icon to move the view in the desired direction.
- To shrink or enlarge the view, click a size factor from the drop-down list box between the magnifying glass icons. Click the - or + magnifying glass icon to modify the size of the view.
- To view a segment of the graph, point to the box to the right of the + magnifying glass icon. Use the red outline in this box to encompass the appropriate segment of the graph. Click the newly outlined area to view the segment.
- To return to the original view, click the magnifying glass icon labeled 1.

System Connectivity view

The System Connectivity view displays the status of system connections in a cluster. Use this view to monitor the system links and disk group heartbeats.

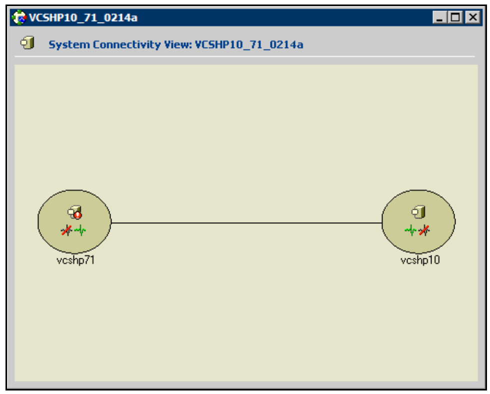

VCS monitors systems and their services over a private network. The systems communicate via heartbeats over an additional private network, which enables them to recognize which systems are active members of the cluster, which are joining or leaving the cluster, and which have failed.

VCS protects against network failure by requiring that all systems be connected by two or more communication channels. When a system is down to a single heartbeat connection, VCS can no longer discriminate between the loss of a system and the loss of a network connection. This situation is referred to as jeopardy.

Point to a system icon to display a ScreenTip on the links and disk group heartbeats. If a system in the cluster is experiencing a problem connecting to other systems, the system icon changes its appearance to indicate the link is down. In this situation, a jeopardy warning may appear in the ScreenTip for this system.

To access the System Connectivity view

1. From Cluster Explorer, click a cluster in the configuration tree.
2. In the view panel, click the **System Connectivity** tab.

Remote Cluster Status view

Note: This view requires the VCS Global Cluster Option.

The Remote Cluster Status View provides an overview of the clusters and global groups in a global cluster environment. Use this view to view the name, address, and status of a cluster, and the type (Icmp or IcmpS) and state of a heartbeat.

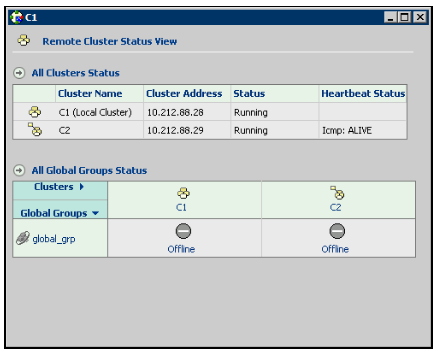

This view enables you to declare a remote cluster fault as a disaster, disconnect, or outage. Point to a table cell to view information about the VCS object.

To access the Remote Cluster Status view

1 From Cluster Explorer, click a cluster in the configuration tree.

2 In the view panel, click the **Remote Cluster Status** tab.

Accessing additional features of the Java Console

Use Cluster Manager to access the Template View, System Manager, User Manager, Command Center, Configuration Wizard, Notifier Resource Configuration Wizard, Remote Group Resource Configuration Wizard, Query Module, and Logs.

Template view

The Template View displays the service group templates available in VCS. Templates are predefined service groups that define the resources, resource attributes, and dependencies within the service group. Use this view to add service groups to the cluster configuration, and copy the resources within a service group template to existing service groups.

In this window, the left pane displays the templates available on the system to which Cluster Manager is connected. The right pane displays the selected template's resource dependency graph.

Template files conform to the VCS configuration language and contain the extension .tf. These files reside in the VCS configuration directory.

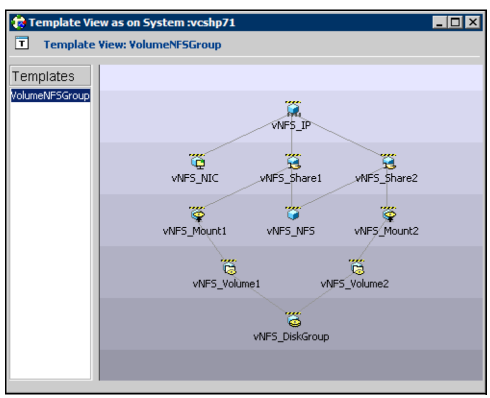

To access the template view

From Cluster Explorer, click **Templates** on the **Tools** menu.

System Manager

Use System Manager to add and remove systems in a service group's system list.

A priority number (starting with 0) is assigned to indicate the order of systems on which the service group will start in case of a failover. If necessary, double-click the entry in the **Priority** column to enter a new value. Select the **Startup** check box to add the systems to the service groups AutoStartList attribute. This enables the service group to automatically come online on a system every time HAD is started.

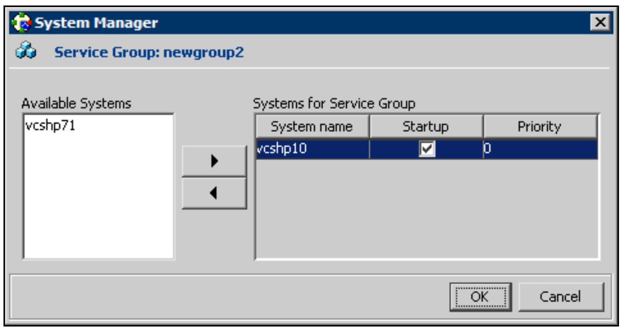

To access system Manager

From Cluster Explorer, click the service group in the configuration tree, and click **System Manager** on the **Tools** menu.

or

In the **Service Groups** tab of the Cluster Explorer configuration tree, click a service group, and click **Manage systems for a Service Group** on the toolbar.

User Manager

User Manager enables you to add and delete user profiles and to change user privileges. If VCS is not running in secure mode, User Manager enables you to change user passwords. You must be logged in as Cluster Administrator to access User Manager.

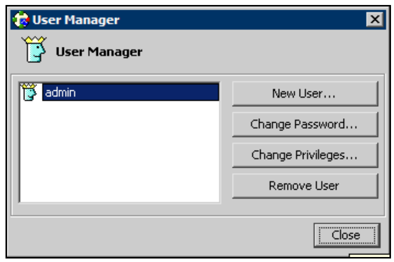

To access user Manager

From Cluster Explorer, click **User Manager** on the **File** menu.

Command Center

Command Center enables you to build and execute VCS commands; most commands that are executed from the command line can also be executed through this window. The left pane of the window displays a **Commands** tree of all VCS operations. The right pane displays a view panel that describes the selected command. The bottom pane displays the commands being executed.

The commands tree is organized into **Configuration** and **Operations** folders. Click the icon to the left of the **Configuration** or **Operations** folder to view its subfolders and command information in the right pane. Point to an entry in the commands tree to display information about the selected command.

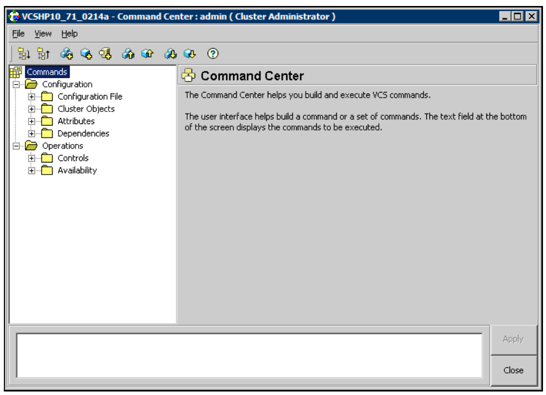

To access Command Center

From Cluster Explorer, click **Command Center** on the **Tools** menu.

or

On the Cluster Explorer toolbar, click **Show Command Center**.

Configuration wizard

Use Configuration Wizard to create and assign service groups to systems in a cluster.

See "Creating service groups with the configuration wizard" on page 197.

To access Configuration Wizard

From Cluster Explorer, click **Configuration Wizard** on the **Tools** menu.

or

On the Cluster Explorer toolbar, click **Launch Configuration Wizard**.

Notifier Resource Configuration wizard

VCS provides a method for notifying an administrator of important events such as a resource or system fault. VCS includes a "notifier" component, which consists of the notifier daemon and the `hanotify` utility. This wizard enables you to configure the notifier component as a resource of type NotifierMngr as part of the ClusterService group.

See "Setting up VCS event notification using the Notifier wizard" on page 227.

To access Notifier Resource Configuration Wizard

From Cluster Explorer, click **Notifier Wizard** on the **Tools** menu.

or

On the Cluster Explorer toolbar, click **Launch Notifier Resource Configuration Wizard**.

Remote Group Resource Configuration Wizard

A RemoteGroup resource enables you to manage or monitor remote service groups from a local cluster. For each service group running in a remote cluster, you can create a corresponding RemoteGroup resource in the local cluster.

See "Adding a RemoteGroup resource from the Java Console" on page 203.

To access Remote Group Resource Configuration Wizard

From Cluster Explorer, click **Remote Group Resource Wizard...** on the **Tools** menu.

or

On the Cluster Explorer toolbar, click **Configure Remote Group Resource Wizard**.

Cluster query

Use Cluster Query to run SQL-like queries from Cluster Explorer. VCS objects that can be queried include service groups, systems, resources, and resource types. Some queries can be customized, including searching for the system's online group count and specific resource attributes.

See "Querying the cluster configuration" on page 226.

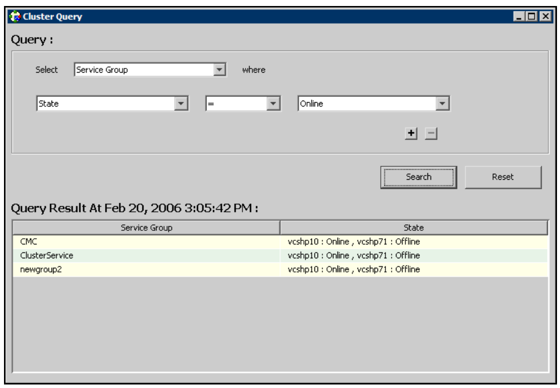

To access the Query dialog box

From Cluster Explorer, click **Query** on the **Tools** menu.

or

In the Cluster Explorer toolbar, click **Query**.

Logs

The Logs dialog box displays the log messages generated by the VCS engine, VCS agents, and commands issued from Cluster Manager to the cluster. Use this dialog box to monitor and take actions on alerts on faulted global clusters and failed service group failover attempts.

Note: To ensure the time stamps for engine log messages are accurate, make sure to set the time zone of the system running the Java Console to the same time zone as the system running the VCS engine.

- Click the **VCS Logs** tab to view the log type, time, and details of an event. Each message presents an icon in the first column of the table to indicate the message type. Use this window to customize the display of messages by setting filter criteria.

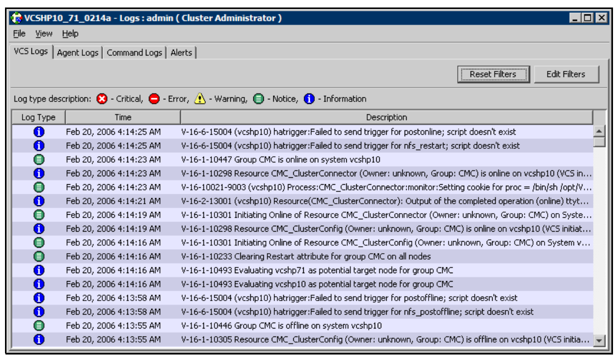

- Click the **Agent Logs** tab to display logs according to system, resource type, and resource filter criteria. Use this tab to view the log type, time, and details of an agent event.
- Click the **Command Logs** tab to view the status (success or failure), time, command ID, and details of a command. The Command Log only displays commands issued in the current session.
- Click the **Alerts** tab to view situations that may require administrative action. Alerts are generated when a local group cannot fail over to any system in the local cluster, a global group cannot fail over, or a cluster fault takes place. A current alert will also appear as a pop-up window when you log on to a cluster through the console.

To access the Logs dialog box

From Cluster Explorer, click **Logs** on the **View** menu.

or

On the Cluster Explorer toolbar, click **Show the Logs**.

Server and user credentials

If VCS is running in secure mode, you can view server and user credentials used to connect to the cluster from Cluster Explorer.

To view user credentials

From Cluster Explorer, click **User Credentials** on the **View** menu.

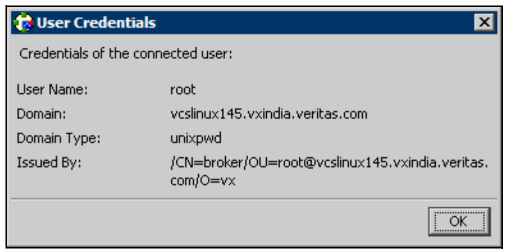

To view server credentials

From Cluster Explorer, click **Server Credentials** on the **View** menu.

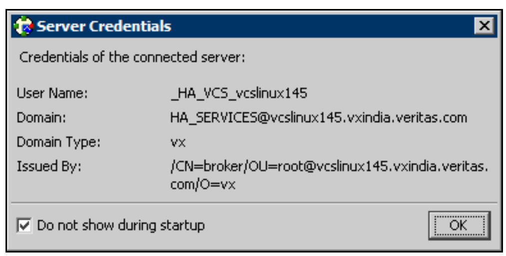

Administering Cluster Monitor

Use the Java Console to administer a cluster or simulated cluster by adding or reconfiguring a cluster panel in Cluster Monitor. To activate the connection of the procedures, log on to the cluster after completing the final step.

Configuring a new cluster panel

You must add a cluster panel for each cluster that you wish to connect to using the Java GUI.

To configure a new cluster panel

1 From Cluster Monitor, click **New Cluster** on the **File** menu. For simulated clusters, click **New Simulator** on the **File** menu.
or
Click **New Cluster** on the Cluster Monitor toolbar.

2 Enter the details to connect to the cluster:

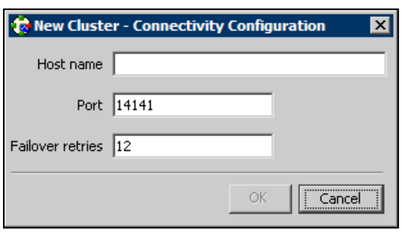

- Enter the host name or IP address of a system in the cluster.
- If necessary, change the default port number of 14141; VCS Simulator uses a default port number of 14153. Note that you must use a different port to connect to each Simulator instance, even if these instances are running on the same system.
- Enter the number of failover retries. VCS sets the default failover retries number to 12.
- For simulated clusters, click the platform for the configuration.
- Click **OK**. An inactive panel appears in Cluster Monitor.

Modifying a cluster panel configuration

Modify a cluster panel to point to another cluster, to change the port number, or the number of failover retries.

1. If Cluster Monitor is in the default expanded state, proceed to step 2. If Cluster Monitor is in the collapsed state:
   On the **View** menu, click **Expand**.
   or
   On the **View** menu, click **Stop** when an active panel appears as the view panel.
2. Right-click the cluster panel. If the panel is inactive, proceed to step 4.
3. On the menu, click **Logout**. The cluster panel becomes inactive.
4. Right-click the inactive panel, and click **Configure...**
5. Edit the details to connect to the cluster:

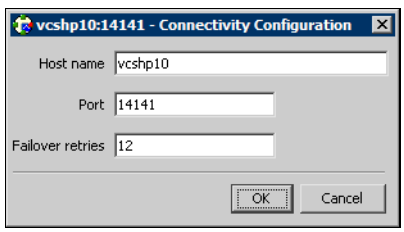

   - Enter the host name or IP address of any system in the cluster.
   - Enter the port number and the number of failover retries. VCS sets the default port number to 14141 and failover retries number to 12; VCS Simulator uses a default port number of 14153.
   - For simulated panels, click the platform for the configuration.
   - Click **OK**.

Logging on to and off of a cluster

After you add or configure a cluster panel in Cluster Monitor, log on to a cluster to access Cluster Explorer. Use Cluster Monitor to log off a cluster when you have completed administering the cluster.

Logging on to a cluster

1 If Cluster Monitor is in the default expanded state, proceed to step 2. If Cluster Monitor is in the collapsed state:
On the **View** menu, click **Expand**.
or
On the **View** menu, click **Stop** when an active panel appears as the view panel.

2 Click the panel that represents the cluster you want to log on to.
or
If the appropriate panel is highlighted, click **Login** on the **File** menu.

3 Enter the information for the user:
If the cluster is not running in secure mode:
- Enter the VCS user name and password.
- Click **OK**.

If the cluster is running in secure mode:

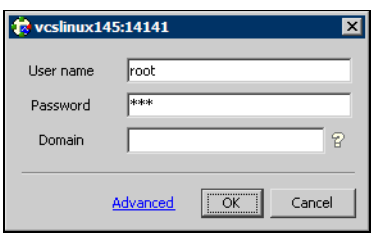

- Enter the credentials of a native user.
You can use nis or nis+ accounts or accounts set up on the local system. If you do not enter the name of the domain, VCS assumes the domain is the local system.
If the user does not have root privileges on the system, VCS assigns guest privileges to the user. To override these privileges, add the domain user to the VCS administrators' list.
See "Administering user profiles" on page 173.

- The Java Console connects to the cluster using the authentication broker and the domain type provided by the engine. To change the authentication broker or the domain type, click **Advanced**.
  See “About security services” on page 29.

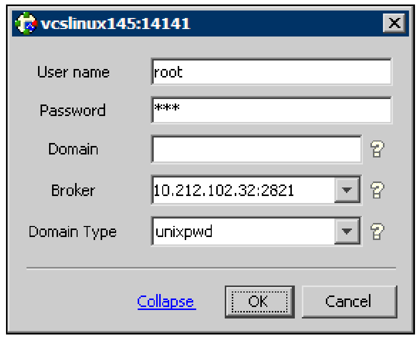

  Select a new broker and domain type, as required.
- Click **OK.**
- The Server Credentials dialog box displays the credentials of the cluster service to which the console is connected.

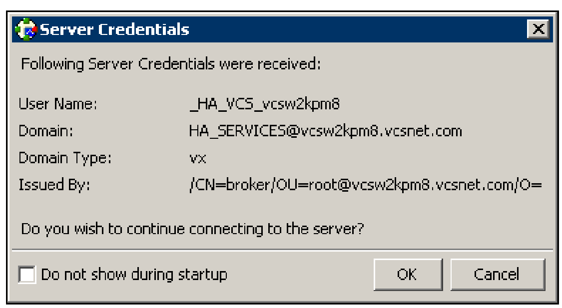

  To disable this dialog box from being displayed every time you connect to the cluster, select the **Do not show during startup** check box
- Click **OK** to connect to the cluster.

The animated display shows various objects, such as service groups and resources, being transferred from the server to the console.

Cluster Explorer is launched automatically upon initial logon, and the icons in the cluster panel change color to indicate an active panel.

Logging off of a cluster

1 If Cluster Monitor is in the default expanded state, proceed to step 2. If Cluster Monitor is in the collapsed state:

   On the **View** menu, click **Expand**.

   or

   On the **View** menu, click **Stop** when an active panel appears as the view panel.

2 Right-click the active panel, and click **Logout**.

   or

   If the appropriate panel is highlighted, click **Logout** on the **File** menu.

   Cluster Explorer closes and the Cluster Monitor panel becomes inactive. You may be prompted to save the configuration if any commands were executed on the cluster.

To log off from Cluster Explorer

Click **Log Out** on the **File** menu.

Administering user profiles

The Java Console enables a user with Cluster Administrator privileges to add, modify, and delete user profiles. The icon next to each user name in the User Manager dialog box indicates privileges for each user. Administrator and Operator privileges are separated into the cluster and group levels.

See "About VCS user privileges and roles" on page 64.

Adding a user

1 From Cluster Explorer, click **User Manager** on the **File** menu.

2 In the User Manager dialog box, click **New User**.

3 In the Add User dialog box:

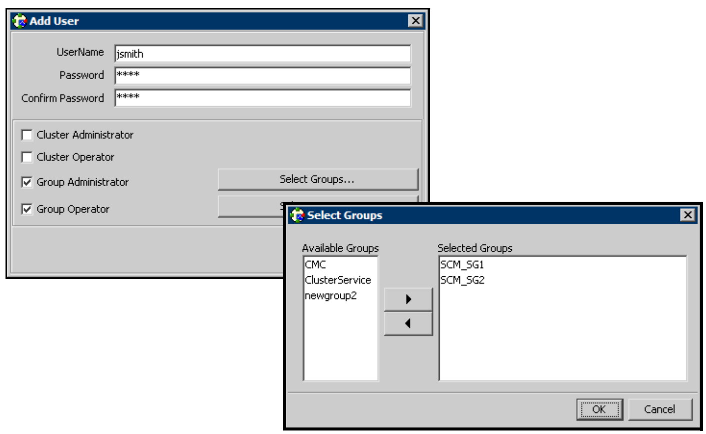

- Enter the name of the user.
- If the cluster is not running in secure mode, enter a password for the user and confirm it.
- Select the appropriate check boxes to grant privileges to the user. To grant Group Administrator or Group Operator privileges, proceed to step the next step. Otherwise, proceed to the last step.
- Click **Select Groups...**
- Click the groups for which you want to grant privileges to the user and click the right arrow to move the groups to the **Selected Groups** box.
- Click **OK** to exit the Select Group dialog box, then click **OK** again to exit the Add User dialog box.

4 Click **Close**.

Deleting a user

1. From Cluster Explorer, click **User Manager** on the **File** menu.
2. In the User Manager dialog box, click the user name.
3. Click **Remove User**.
4. Click **Yes**.
5. Click **Close**.

Changing a user password

A user with Administrator, Operator, or Guest privileges can change his or her own password. You must be logged on as Cluster Administrator to access User Manager. Before changing the password, make sure the configuration is in the read-write mode. Cluster administrators can change the configuration to the read-write mode.

Note: This module is not available if the cluster is running in secure mode.

To change a password as an administrator

1. From Cluster Explorer, click **User Manager** on the **File** menu.
2. Click the user name.
3. Click **Change Password**.
4. In the Change Password dialog box:
   - Enter the new password.
   - Re-enter the password in the **Confirm Password** field.
   - Click **OK**.
5. Click **Close**.

To change a password as an operator or guest

1. From Cluster Explorer, click **Change Password** on the **File** menu.
2. In the Change Password dialog box:
   - Enter the new password.
   - Reenter the password in the **Confirm Password** field.
   - Click **OK**.
3. Click **Close**.

Changing a user privilege

1. From Cluster Explorer, click **User Manager** on the **File** menu.
2. Click the user name.
3. Click **Change Privileges** and enter the details for user privileges:

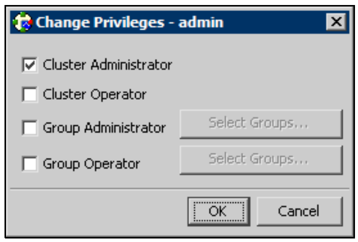

- Select the appropriate check boxes to grant privileges to the user. To grant Group Administrator or Group Operator privileges, proceed to the next step. Otherwise, proceed to the last step.
- Click **Select Groups**.
- Click the groups for which you want to grant privileges to the user, then click the right arrow to move the groups to the **Selected Groups** box.
- Click **OK** in the Change Privileges dialog box, then click **Close** in the User Manager dialog box.

Assigning privileges for OS user groups in secure clusters

In secure clusters, you can assign privileges to native users at an operating system (OS) user group level. Assigning VCS privileges to an OS user group involves adding the user group in one (or more) of the following attributes:

- AdministratorGroups—for a cluster or for a service group.
- OperatorGroups—for a cluster or for a service group.
- Guests

See "User privileges for OS user groups in secure clusters" on page 68.

To assign privileges to an OS user group

1. From Cluster Explorer configuration tree, select the cluster to assign privileges for the cluster or a service group to assign privileges for specific service groups.
2. From the view panel, click the **Properties** tab.
3. From the list of key attributes, click the edit icon against **AdministratorGroups**, **OperatorGroups**, or **Guests** attributes.
4. In the Edit Attribute dialog box:
   - Use the **+** button to add an element.
   - Click the newly added element and enter the name of the user group in the format *group@domain*.
   - Click **OK**.

Administering service groups

Use the Java Console to administer service groups in the cluster. Use the console to add and delete, bring online and take offline, freeze and unfreeze, link and unlink, enable and disable, autoenable, switch, and flush service groups. You can also modify the system list for a service group.

Adding a service group

The Java Console provides several ways to add a service group to the systems in a cluster. Use Cluster Explorer, Command Center, or the Template View to perform this task.

To add a service group from Cluster Explorer

1 On the **Edit** menu, click **Add**, and click **Service Group**.
or
Click **Add Service Group** in the Cluster Explorer toolbar.

2 Enter the details of the service group:

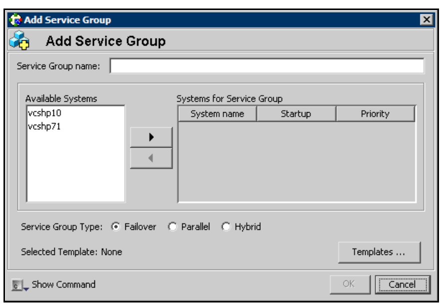

- Enter the name of the service group.
- In the **Available Systems** box, click the systems on which the service group will be added.
- Click the right arrow to move the selected systems to the **Systems for Service Group** box. The priority number (starting with 0) is automatically assigned to indicate the order of systems on which the service group will start in case of a failover. If necessary, double-click the entry in the **Priority** column to enter a new value.
  Select the **Startup** check box to add the systems to the service groups AutoStartList attribute. This enables the service group to automatically come online on a system every time HAD is started.

- Click the appropriate service group type. A failover service group runs on only one system at a time; a parallel service group runs concurrently on multiple systems.
- To add a new service group based on a template, click **Templates...** Otherwise, proceed to step 2g. (Alternative method to add a new service group based on a template: From Cluster Explorer, click **Templates** on the **Tools** menu. Right-click the Template View panel, and click **Add as Service Group** from the menu.)
- Click the appropriate template name, then click **OK**.
- Click **Show Command** in the bottom left corner if you want to view the command associated with the service group. Click **Hide Command** to close the view of the command.
- Click **OK**.

To add a service group from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Cluster Objects** > **Add Service Group**.

or

Click **Add service group** in the Command Center toolbar.

2 Enter the name of the service group.

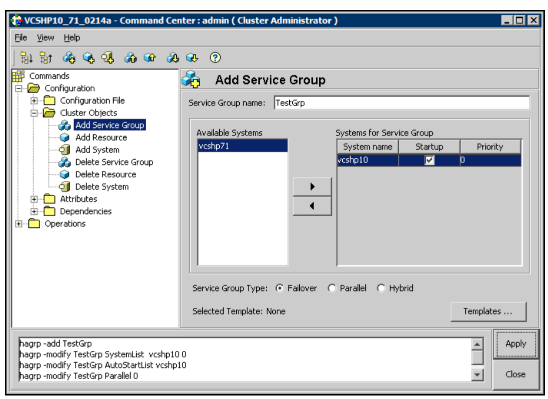

3 In the **Available Systems** box, click the systems on which the service group will be added.

4 Click the right arrow to move the selected systems to the **Systems for Service Group** box. The priority number (starting with 0) is automatically assigned to indicate the order of systems on which the service group will start in case of a failover. If necessary, double-click the entry in the **Priority** column to enter a new value.

Select the **Startup** check box to add the systems to the service groups AutoStartList attribute. This enables the service group to automatically come online on a system every time HAD is started.

5 Click the appropriate service group type. A failover service group runs on only one system at a time; a parallel service group runs concurrently on multiple systems.

6 To add a new service group based on a template, click **Templates...** Otherwise, proceed to step 9.

7 Click the appropriate template name.

8 Click **OK**.

9 Click **Apply**.

To add a service group from the template view

1 From Cluster Explorer, click **Templates...** on the **Tools** menu.

2 Right-click the Template View panel, and click **Add as Service Group** from the pop-up menu. This adds the service group template to the cluster configuration file without associating it to a particular system.

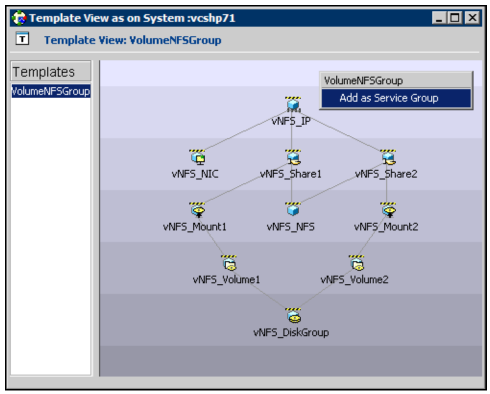

3 Use System Manager to add the service group to systems in the cluster. See "System Manager" on page 160.

Deleting a service group

Delete a service group from Cluster Explorer or Command Center.

Note: You cannot delete service groups with dependencies. To delete a linked service group, you must first delete the link.

To delete a service group 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 **Delete** from the menu.

3 Click **Yes**.

To delete a service group from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Cluster Objects** > **Delete Service Group**.

2 Click the service group.

3 Click **Apply**.

Bringing a service group online

To bring a service group online from the Cluster Explorer configuration tree

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 the appropriate system from the menu. Click **Any System** if you do not need to specify a system.

To bring a service group online from the Cluster Explorer toolbar

1. Click **Online Service Group** on the Cluster Explorer toolbar.
2. Specify the details for the service group:

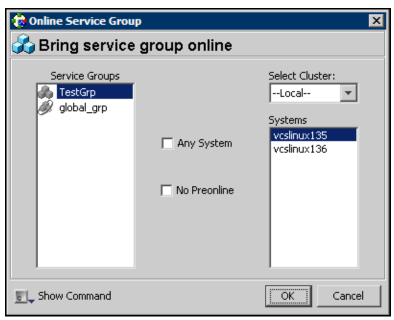

- Click the service group.
- For global groups, select the cluster in which to bring the group online.
- Click the system on which to bring the group online, or select the **Any System** check box.
- Select the **No Preonline** check box to bring the service group online without invoking the preonline trigger.
- Click **Show Command** in the bottom left corner to view the command associated with the service group. Click **Hide Command** to close the view of the command.
- Click **OK**.

To bring a service group online from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **Online Service Group**.
   or
   Click **Bring service group online** in the Command Center toolbar.
2. Click the service group.
3. For global groups, select the cluster in which to bring the group online.
4. Click the system on which to bring the group online, or select the **Any System** check box.
5. Click **Apply**.

Taking a service group offline

To take a service group offline from Cluster Explorer configuration tree

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 Group**s tab, and right-click the service group icon in the view panel.
2. Click **Offline**, and click the appropriate system from the menu. Click **All Systems** to take the group offline on all systems.

To take a service group offline from the Cluster Explorer toolbar

1. Click **Offline Service Group** in the Cluster Explorer toolbar.
2. Enter the details of the service group:

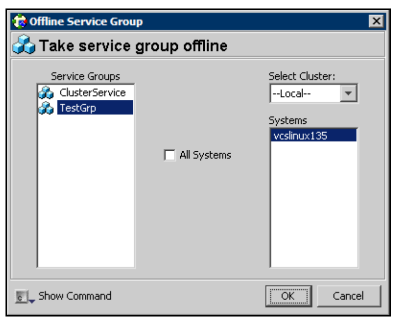

- Click the service group.
- For global groups, select the cluster in which to take the group offline.
- Click the system on which to take the group offline, or click **All Systems**.
- Click **Show Command** in the bottom left corner if you want to view the command associated with the service group. Click **Hide Command** to close the view of the command.
- Click **OK**.

To take a service group offline from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **Offline Service Group**.
   or
   Click **Take service group offline** in the Command Center toolbar.
2. Click the service group.
3. For global groups, select the cluster in which to take the group offline.
4. Click the system on which to take the group offline, or click the **All Systems** check box.
5. Click **Apply**.

Switching a service group

The process of switching a service group involves taking it offline on its current system and bringing it online on another system.

To switch a service group from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the service group.
   or
   Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
2. Click **Switch To**, and click the appropriate system from the menu.

To switch a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **Switch Service Group**.
2. Click the service group.
3. For global groups, select the cluster in which to switch the service group.
4. Click the system on which to bring the group online, or select the **Any System** check box.
5. Click **Apply**.

Freezing a service group

Freeze a service group to prevent it from failing over to another system. The freezing process stops all online and offline procedures on the service group.

To freeze a service group from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the service group.
   or
   Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
2. Click **Freeze**, and click **Temporary** or **Persistent** from the menu. The persistent option maintains the frozen state after a reboot if you save this change to the configuration.

To freeze a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Freeze Service Group**.
2. Click the service group.
3. Select the **persistent** check box if necessary. The persistent option maintains the frozen state after a reboot if you save this change to the configuration.
4. Click **Apply**.

Unfreezing a service group

Unfreeze a frozen service group to perform online or offline operations on the service group.

To unfreeze a service group from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the service group.
   or
   Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
2. Click **Unfreeze**.

To unfreeze a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Unfreeze Service Group**.
2. Click the service group.
3. Click **Apply**.

Enabling a service group

Enable a service group before bringing it online. A service group that was manually disabled during a maintenance procedure on a system may need to be brought online after the procedure is completed.

To enable a service group from Cluster Explorer

1 In the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.

2 Click **Enable**, and click the appropriate system from the menu. Click **All Systems** to enable the group on all systems.

To enable a service group from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Enable Service Group**.

2 Click the service group.

3 Select the **Per System** check box to enable the group on a specific system instead of all systems.

4 Click **Apply**.

Disabling a service group

Disable a service group to prevent it from coming online. This process temporarily stops VCS from monitoring a service group on a system undergoing maintenance operations.

To disable a service group from Cluster Explorer

1 In the **Service Groups** tab of the configuration tree, right-click the service group.
  or
  Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.

2 Click **Disable**, and click the appropriate system in the menu. Click **All Systems** to disable the group on all systems.

To disable a service group from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Disable Service Group**.

2 Click the service group.

3 Select the **Per System** check box to disable the group on a specific system instead of all systems.

4 Click **Apply**.

Autoenabling a service group

A service group is autodisabled until VCS probes all resources and checks that they are ready to come online. Autoenable a service group in situations where the VCS engine is not running on one of the systems in the cluster, and you must override the disabled state of the service group to enable the group on another system in the cluster.

To autoenable a service group from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the service group.
   or
   Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
2. Click **Autoenable**, and click the appropriate system from the menu.

To autoenable a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Autoenable Service Group**.
2. Click the service group.
3. Click the system on which to autoenable the group.
4. Click **Apply**.

Flushing a service group

As a service group is brought online or taken offline, the resources within the group are brought online and taken offline. If the online or offline operation hangs on a particular resource, flush the service group to halt the operation on the resources waiting to go online or offline. Flushing a service group typically leaves the cluster in a partial state. After completing this process, resolve the issue with the particular resource (if necessary) and proceed with starting or stopping the service group.

To flush a service group from Cluster Explorer

1 In the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.

2 Click **Flush**, and click the appropriate system from the menu.

To flush a service group from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Flush Service Group**.

2 Click the service group.

3 Click the system on which to flush the service group.

4 Click **Apply**.

Linking service groups

To link a service group from Cluster Explorer

1. Click a cluster in the configuration tree.
2. In the view panel, click the **Service Groups** tab. This opens the service group dependency graph. To link a parent group with a child group:
   - Click **Link**.
   - Click the parent group.
   - Move the mouse toward the child group. The yellow line “snaps” to the child group. If necessary, press Esc on the keyboard to delete the line between the parent and the pointer before it snaps to the child.
   - Click the child group.
   - In the Link Service Groups dialog box, click the group relationship and dependency type.
     See “About service group dependencies” on page 492.

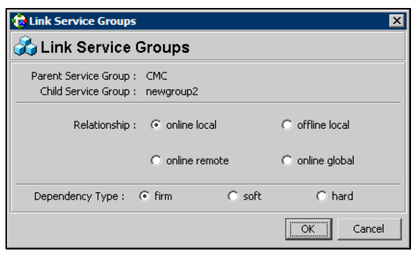

   - Click **OK**.

You can also link the service groups by performing steps 1 and 2, right-clicking the parent group, and clicking **Link** from the menu. In the dialog box, click the child group, relationship, dependency type, and click **OK**.

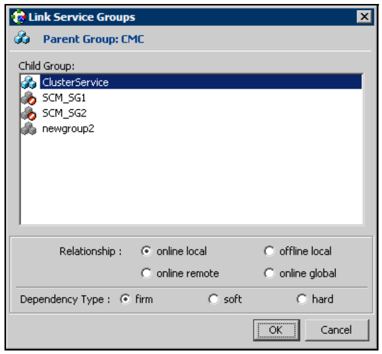

To link a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Configuration** > **Dependencies** > **Link Service Groups**.
2. Click the parent resource group in the **Service Groups** box. After selecting the parent group, the potential groups that can serve as child groups are displayed in the **Child Service Groups** box.
3. Click a child service group.
4. Click the group relationship and dependency type.
   See "About service group dependencies" on page 492.
5. Click **Apply**.

Unlinking service groups

To delete a service group dependency from Cluster Explorer

1. Click a cluster in the configuration tree.
2. In the view panel, click the **Service Groups** tab.
3. In the Service Group view, right-click the link between the service groups.
4. Click **Unlink** from the menu.

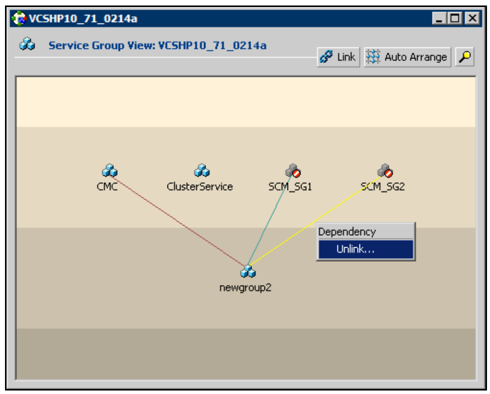

5. Click **Yes**.

To delete a service group dependency from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Configuration** > **Dependencies** > **Unlink Service Groups**.
2. Click the parent resource group in the **Service Groups** box. After selecting the parent group, the corresponding child groups are displayed in the **Child Service Groups** box.
3. Click the child service group.
4. Click **Apply**.

Managing systems for a service group

From Cluster Explorer, use System Manager to add and remove systems in a service group's system list.

To add a system to the service group's system list

1. In the System Manager dialog box, click the system in the **Available Systems** box.

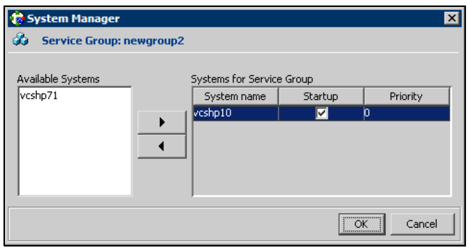

2. Click the right arrow to move the available system to the **Systems for Service Group** table.
3. Select the **Startup** check box to add the systems to the service groups AutoStartList attribute. This enables the service group to automatically come online on a system every time HAD is started.
4. The priority number (starting with 0) is assigned to indicate the order of systems on which the service group will start in case of a failover. If necessary, double-click the entry in the **Priority** column to enter a new value.
5. Click **OK**.

To remove a system from the service group's system list

1. In the System Manager dialog box, click the system in the **Systems for Service Group** table.
2. Click the left arrow to move the system to the **Available Systems** box.
3. Click **OK**.

Creating service groups with the configuration wizard

This section describes how to create service groups using the configuration wizard.

Note: VCS also provides wizards to create service groups for applications and NFS shares. See the chapter "Configuring Application and NFS Service Groups" on page 291 for more information about these wizards.

To create a service group using the configuration wizard

1 Open the Configuration Wizard. From Cluster Explorer, click **Configuration Wizard** on the **Tools** menu.

2 Read the information on the Welcome dialog box and click **Next**.

3 Specify the name and target systems for the service group:

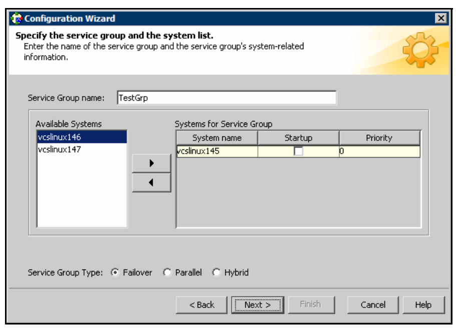

- Enter the name of the group.
- Click the target systems in the **Available Systems** box.
- Click the right arrow to move the systems to the **Systems for Service Group** table. To remove a system from the table, click the system and click the left arrow.
- Select the **Startup** check box to add the systems to the service groups AutoStartList attribute. This enables the service group to automatically come online on a system every time HAD is started.
- The priority number (starting with 0) is automatically assigned to indicate the order of systems on which the service group will start in case of a failover. If necessary, double-click the entry in the **Priority** column to enter a new value.
- Click the service group type.

- Click **Next**.

4 Click **Next** again to configure the service group with a template and proceed to step 7. Click **Finish** to add an empty service group to the selected cluster systems and configure it at a later time.

5 Click the template on which to base the new service group. The Templates box lists the templates available on the system to which Cluster Manager is connected. The resource dependency graph of the templates, the number of resources, and the resource types are also displayed. Click **Next**.

6 If a window notifies you that the name of the service group or resource within the service group is already in use, proceed to step 9. Otherwise, proceed to step 10.

7 Click **Next** to apply all of the new names listed in the table to resolve the name clash.
or
Modify the clashing names by entering text in the field next to the **Apply** button, clicking the location of the text for each name from the **Correction** drop-down list box, clicking **Apply**, and clicking **Next**.

8 Click **Next** to create the service group. A progress indicator displays the status.

9 After the service group is successfully created, click **Next** to edit attributes using the wizard. Click **Finish** to edit attributes at a later time using Cluster Explorer.

10 Review the attributes associated with the resources of the service group. If necessary, proceed to step 11 to modify the default values of the attributes. Otherwise, proceed to step 12 to accept the default values and complete the configuration.

11 Modify the values of the attributes (if necessary).

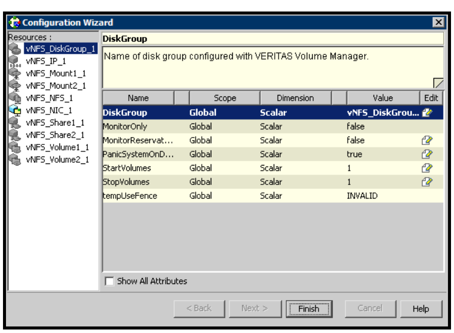

- Click the resource.

- Click the attribute to be modified.
- Click the **Edit** icon at the end of the table row.
- In the Edit Attribute dialog box, enter the attribute values.
- Click **OK**.
- Repeat the procedure for each resource and attribute.

12 Click **Finish**.

Administering resources

Use the Java Console to administer resources in the cluster. Use the console to add and delete, bring online and take offline, probe, enable and disable, clear, and link and unlink resources. You can also import resource types to the configuration.

Adding a resource

The Java Console provides several ways to add a resource to a service group. Use Cluster Explorer or Command Center to perform this task.

To add a resource from Cluster Explorer

1 In the **Service Groups** tab of the Cluster Explorer configuration tree, click a service group to which the resource will be added.

2 On the **Edit** menu, click **Add**, and click **Resource**.
or
Click **Add Resource** in the Cluster Explorer toolbar.

3 Enter the details of the resource:

- Enter the name of the resource.

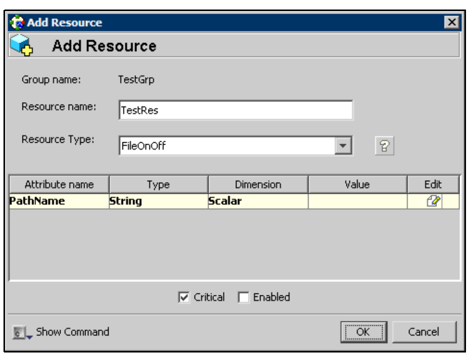

- Click the resource type.
- Edit resource attributes according to your configuration. The Java Console also enables you to edit attributes after adding the resource.
- Select the **Critical** and **Enabled** check boxes, if applicable. The **Critical** option is selected by default.
  A critical resource indicates the service group is faulted when the resource, or any resource it depends on, faults. An enabled resource indicates agents monitor the resource; you must specify the values of mandatory attributes before enabling a resource. If a resource is

created dynamically while VCS is running, you must enable the resource before VCS monitors it. VCS will not bring a disabled resource nor its children online, even if the children are enabled.

- Click **Show Command** in the bottom left corner to view the command associated with the resource. Click **Hide Command** to close the view of the command.
- Click **OK**.

To add a resource from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Cluster Objects** > **Add Resource**.
or
Click **Add resource** in the Command Center toolbar.

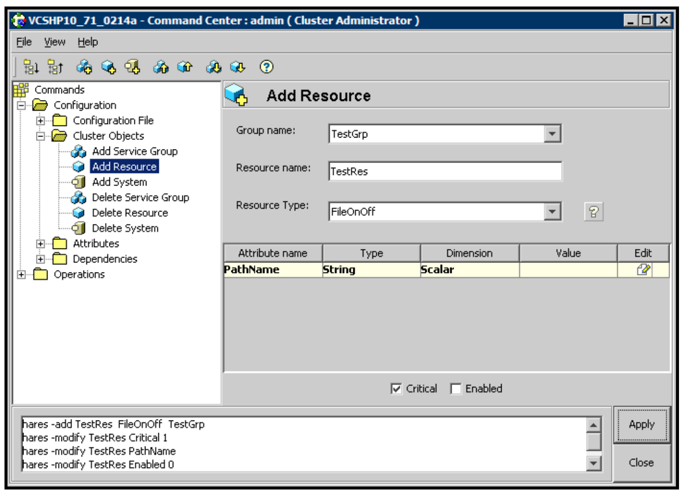

2 Select the service group to contain the resource.

3 Enter the name of the resource.

4 Click the resource type.

5 Edit resource attributes according to your configuration. The Java Console also enables you to edit attributes after adding the resource.

6 Select the **Critical** and **Enabled** check boxes, if applicable. The **Critical** option is selected by default.
A critical resource indicates the service group is faulted when the resource, or any resource it depends on, faults. An enabled resource indicates agents monitor the resource; you must specify the values of mandatory attributes

before enabling a resource. If a resource is created dynamically while VCS is running, you must enable the resource before VCS monitors it. VCS will not bring a disabled resource nor its children online, even if the children are enabled.

7 Click **Apply**.

To add a resource from the Template view

1 From Cluster Explorer, click **Templates...** on the **Tools** menu.

2 In the left pane of the Template View, click the template from which to add resources to your configuration.

3 In the resource graph, right-click the resource to be added to your configuration.

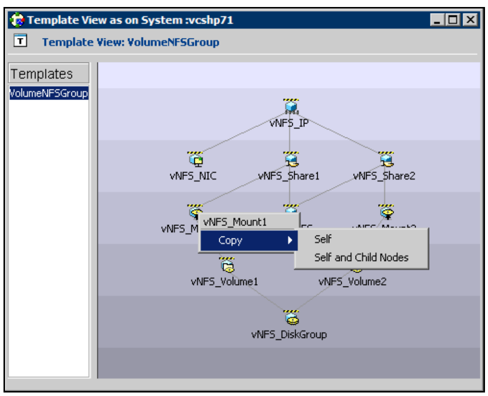

4 Click **Copy**, and click **Self** from the menu to copy the resource. Click **Copy**, and click **Self and Child Nodes** from the menu to copy the resource with its dependent resources.

5 In the **Service Groups** tab of the Cluster Explorer configuration tree, click the service group to which to add the resources.

6 In the Cluster Explorer view panel, click the **Resources** tab.

7 Right-click the Resource view panel and click **Paste** from the menu. After the resources are added to the service group, edit the attributes to configure the resources.

Adding a RemoteGroup resource from the Java Console

A RemoteGroup resource is typically useful in scenarios where resources configured in a local service group are dependant on the state of a remote service group. For example, a web-server application running in a local cluster could be dependant on a database application running in a remote cluster.

A RemoteGroup resource monitors the state of a remote service group in a local cluster. Once you have added the RemoteGroup resource to a local service group, you can link the resource to the existing resources of the service group.

You must have administrative privileges to configure RemoteGroup resources.

See “Configuring the RemoteGroup agent” on page 242.

To add a RemoteGroup resource

1. On the **Tools** menu, click **Add Remote Group Resource...**

   or

   Click **Configure Remote Group Resource Wizard** in the Cluster Explorer toolbar.

2. Read the information on the Welcome dialog box and click **Next**.

3. In the Remote Group Resource Name dialog box, specify the name of the resource and the service group to which the resource will be added. Click **Next**.

4. In the Remote Cluster Information dialog box:
   - Specify the name or IP address of a node in the remote cluster.
   - Specify the port on the remote node on which the resource will communicate.
   - Specify a username for the remote cluster.
   - Specify a password for the user.
   - Select the check box if you wish to specify advance options to connect to a secure cluster. Otherwise, proceed to the last step.
     - Specify the domain of which the node is a part.
     - Select a domain type.
     - Specify the authentication broker and port.
   - Click **Next**.

5 In the Remote Group Resource Details dialog box:

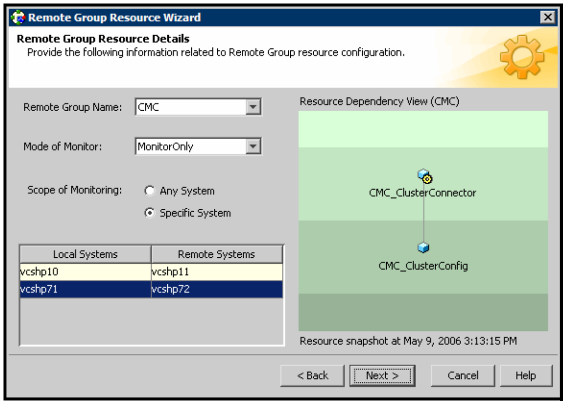

- Select a group you wish to monitor.
- Select the mode of monitoring.
  - Choose the **MonitorOnly** option to monitor the remote service group. You will not be able to perform online or offline operations on the remote group.
  - Choose the **OnlineOnly** option to monitor the remote service group and bring the remote group online from the local cluster.
  - Choose the **OnOff** option to monitor the remote service group, bring the remote group online, and take it offline from the local cluster.
- Specify whether the RemoteGroup resource should monitor the state of the remote group on a specific system or any system in the remote cluster.
  - Choose the **Any System** option to enable the RemoteGroup resource to monitor the state of the remote service group irrespective of the system on which it is online.
  - Choose the **Specific System** option to enable the RemoteGroup resource to monitor the state of the remote group on a specific system in the remote cluster. Both service groups must be configured on the same number of systems.

    This option provides one-to-one mapping between the local and remote systems. The **Local Systems** list displays the systems on which the RemoteGroup resource is configured. Click the fields under the **Remote Systems** list and select the systems from drop-down list. If the remote group fails over to another system in the remote cluster, the RemoteGroup resource will also fail over to the corresponding system in the local cluster.

- Click **Next**.

6 Review the text in the dialog box and click **Finish** to add the RemoteGroup resource to the specified service group in the local cluster. You must now create dependencies between the RemoteGroup resource and the existing resources of the service group.
See "Linking resources" on page 213.

Deleting a resource

To delete a resource from Cluster Explorer

1 In the **Service Groups** tab of the configuration tree, right-click the resource.
or
Click a service group in the configuration tree, click the **Resources** tab, and right-click the resource icon in the view panel.

2 Click **Delete** from the menu.

3 Click **Yes**.

To delete a resource from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Cluster Objects** > **Delete Resource**.

2 Click the resource.

3 Click **Apply**.

Bringing a resource online

To bring a resource online from Cluster Explorer

1 In the **Service Groups** tab of the configuration tree, right-click the resource.
or
Click a service group in the configuration tree, click the **Resources** tab, and right-click the resource icon in the view panel.

2 Click **Online**, and click the appropriate system from the menu.

To bring a resource online from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **Online Resource**.

2 Click a resource.

3. Click a system on which to bring the resource online.
4. Click **Apply**.

Taking a resource offline

To take a resource offline from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the resource.
   or
   Click a service group in the configuration tree, click the **Resources** tab, and right-click the resource icon in the view panel.
2. Click **Offline**, and click the appropriate system from the menu.

To take a resource offline from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **Offline Resource**.
2. Click a resource.
3. Click a system on which to take the resource offline.
4. If necessary, select the **ignoreparent** check box to take a selected child resource offline, regardless of the state of the parent resource. This option is only available through Command Center.
5. Click **Apply**.

Taking a resource offline and propagating the command

Use the Offline Propagate (OffProp) feature to propagate the offline state of a parent resource. This command signals that resources dependent on the parent resource should also be taken offline.

Use the Offline Propagate (OffProp) “ignoreparent” feature to take a selected resource offline, regardless of the state of the parent resource. This command propagates the offline state of the selected resource to the child resources. The “ignoreparent” option is only available in Command Center.

To take a resource and its child resources offline from Cluster Explorer

1. In the Resources tab of the configuration tree, right-click the resource.
2. Click **Offline Prop**, and click the appropriate system from the menu.

To take a resource and its child resources offline from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **OffProp Resource**.
2. Click the resource.
3. Click the system on which to take the resource, and the child resources, offline.
4. Click **Apply**.

To take child resources offline from Command Center while ignoring the state of the parent resource

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **OffProp Resource**.
2. Click the resource.
3. Click the system on which to take the resource, and the child resources, offline.
4. Select the **ignoreparent** check box.
5. Click **Apply**.

Probing a resource

Probe a resource to check that it is configured and ready to bring online.

To probe a resource from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the resource.
2. Click **Probe**, and click the appropriate system from the menu.

To probe a resource from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Controls** > **Probe Resource**.
2. Click the resource.
3. Click the system on which to probe the resource.
4. Click **Apply**.

Overriding resource type static attributes

You can override some resource type static attributes and assign them resource-specific values. When a static attribute is overridden and the configuration is saved, the main.cf file includes a line in the resource definition for the static attribute and its overridden value.

To override resource type static attribute

1. Right-click the resource in the **Service Groups** tab of the configuration tree or in the **Resources** tab of the view panel.
2. Click **Override Attributes**.
3. Select the attributes to override.
4. Click **OK**.
   The selected attributes appear in the Overridden Attributes table in the Properties view for the resource.
5. To modify the default value of an overridden attribute, click the icon in the **Edit** column of the attribute.

To restore default settings to a type's static attribute

1. Right-click the resource in the **Service Groups** tab of the configuration tree or in the **Resources** tab of the view panel.
2. Click **Remove Attribute Overrides**.
3. Select the overridden attributes to be restored to their default settings.
4. Click **OK**.

Enabling resources in a service group

Enable resources in a service group to bring the disabled resources online. A resource may have been manually disabled to temporarily stop VCS from monitoring the resource. You must specify the values of mandatory attributes before enabling a resource.

To enable an individual resource in a service group

1. From Cluster Explorer, click the **Service Groups** tab of the configuration tree.
2. Right-click a disabled resource in the configuration tree, and click **Enabled** from the menu.

To enable all resources in a service group from Cluster Explorer

1. From Cluster Explorer, click the **Service Groups** tab in the configuration tree.
2. Right-click the service group.
3. Click **Enable Resources**.

To enable all resources in a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Enable Resources for Service Group**.
2. Click the service group.
3. Click **Apply**.

Disabling resources in a service group

Disable resources in a service group to prevent them from coming online. This disabling process is useful when you want VCS to temporarily "ignore" resources (rather than delete them) while the service group is still online.

To disable an individual resource in a service group

1. From Cluster Explorer, click the **Service Groups** tab in the Cluster Explorer configuration tree.
2. Right-click a resource in the configuration tree. An enabled resource will display a check mark next to the **Enabled** option that appears in the menu.
3. Click **Enabled** from the menu to clear this option.

To disable all resources in a service group from Cluster Explorer

1. From Cluster Explorer, click the **Service Groups** tab in the configuration tree.
2. Right-click the service group and click **Disable Resources**.

To disable all resources in a service group from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Disable Resources for Service Group**.
2. Click the service group.
3. Click **Apply**.

Clearing a resource

Clear a resource to remove a fault and make the resource available to go online. A resource fault can occur in a variety of situations, such as a power failure or a faulty configuration.

To clear a resource from Cluster Explorer

1. In the **Service Groups** tab of the configuration tree, right-click the resource.
2. Click **Clear Fault**, and click the system from the menu. Click **Auto** instead of a specific system to clear the fault on all systems where the fault occurred.

To clear a resource from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Clear Resource**.
2. Click the resource. To clear the fault on all systems listed in the **Systems** box, proceed to step 5. To clear the fault on a specific system, proceed to step 3.
3. Select the **Per System** check box.
4. Click the system on which to clear the resource.
5. Click **Apply**.

Linking resources

Use Cluster Explorer or Command Center to link resources in a service group.

To link resources from Cluster Explorer

1. In the configuration tree, click the **Service Groups** tab.
2. Click the service group to which the resources belong.
3. In the view panel, click the **Resources** tab. This opens the resource dependency graph. To link a parent resource with a child resource:
   - Click **Link...**
   - Click the parent resource.
   - Move the mouse towards the child resource. The yellow line “snaps” to the child resource. If necessary, press Esc to delete the line between the parent and the pointer before it snaps to the child.
   - Click the child resource.
   - In the Confirmation dialog box, click **Yes**.
     or
     Right-click the parent resource, and click **Link** from the menu. In the Link Resources dialog box, click the resource that will serve as the child. Click **OK**.

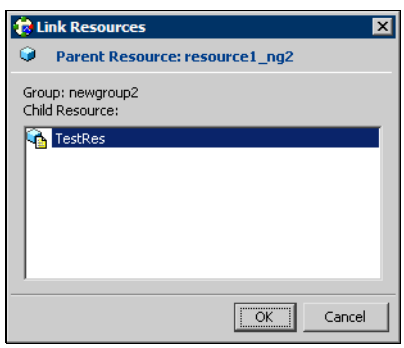

   - Click **OK**.

To link resources from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Configuration** > **Dependencies** > **Link Resources**.
2. Click the service group to contain the linked resources.

3 Click the parent resource in the **Service Group Resources** box. After selecting the parent resource, the potential resources that can serve as child resources are displayed in the **Child Resources** box.

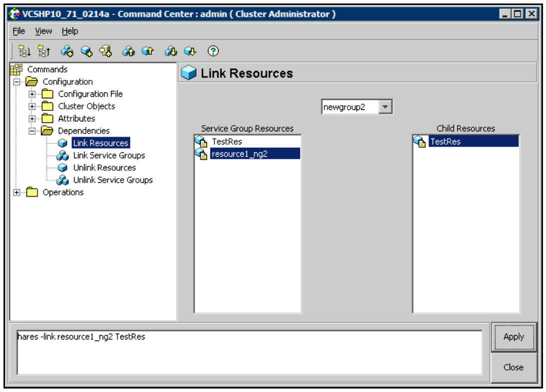

4 Click a child resource.

5 Click **Apply**.

Unlinking resources

Use Cluster Explorer or Command Center to unlink resources in a service group.

To unlink resources from Cluster Explorer

1. From the configuration tree, click the **Service Groups** tab.
2. Click the service group to which the resources belong.
3. In the view panel, click the **Resources** tab.
4. In the Resources View, right-click the link between the resources.
5. Click **Unlink...** from the menu.

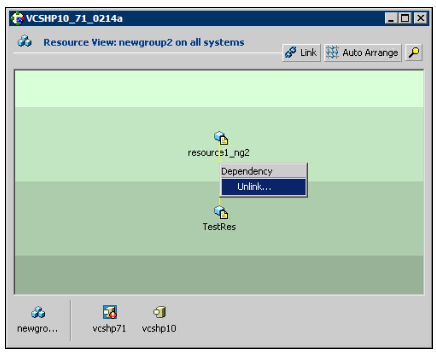

6. In the Question dialog box, click **Yes** to delete the link.

To unlink resources from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Dependencies** > **Unlink Resources**.

2 Click the service group that contains the linked resources.

3 Click the parent resource in the **Service Group Resources** box. After selecting the parent resource, the corresponding child resources are displayed in the **Child Resources** box.

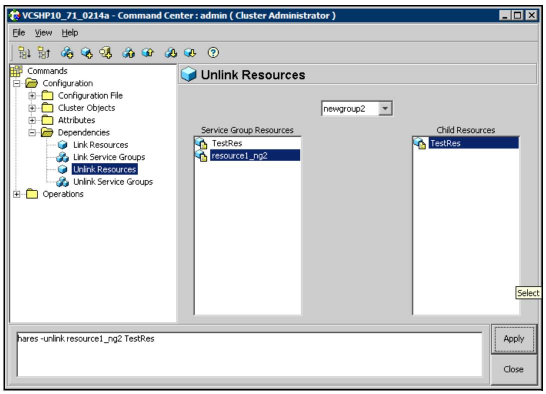

4 Click the child resource.

5 Click **Apply**.

Invoking a resource action

Cluster Explorer enables you to initiate a predefined action script. Some examples of predefined resource actions are splitting and joining disk groups.

To invoke a resource action

1 In the **Service Groups** tab of the configuration tree, right-click the resource.

2 Click **Actions...**

3 Specify the details of the action:

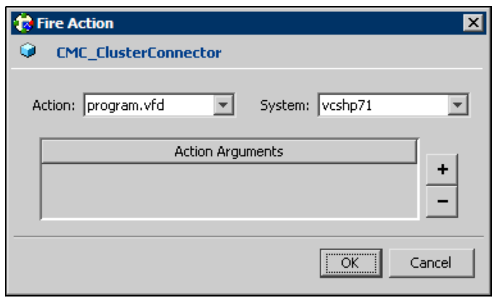

- Click the predefined action to execute.
- Click the system on which to execute the action.
- To add an argument, click the **Add** icon (**+**) and enter the argument. Click the **Delete** icon (**-**) to remove an argument.
- Click **OK**.

Refreshing the ResourceInfo attribute

Refresh the ResourceInfo attribute to view the latest values for that attribute.

To refresh the ResourceInfo attribute

1. In the **Service Groups** tab of the configuration tree, right-click the resource.
2. Click **Refresh ResourceInfo**, and click the system on which to refresh the attribute value.

Clearing the ResourceInfo attribute

Clear the ResourceInfo attribute to reset all the parameters in this attribute.

To clear the parameters of the ResourceInfo attribute

1. In the **Service Groups** tab of the configuration tree, right-click the resource.
2. Click **Clear ResourceInfo**, and click the system on which to reset the attribute value.

Importing resource types

The Java Console enables you to import resource types into your configuration (main.cf). For example, use this procedure to import the types.cf for enterprise agents to your configuration. You cannot import resource types that already exist in your configuration.

To import a resource type from Cluster Explorer

1. On the **File** menu, click **Import Types**.
2. In the Import Types dialog box:
   - Click the file from which to import the resource type. The dialog box displays the files on the system that Cluster Manager is connected to.
   - Click **Import**.

Administering systems

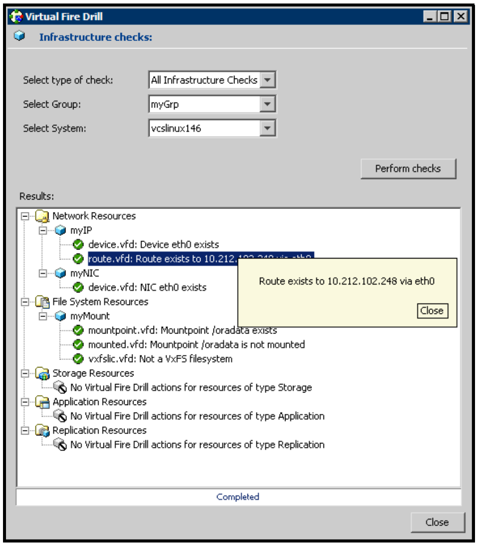

Use the Java Console to administer systems in the cluster. Use the console to add, delete, freeze, and unfreeze systems.

Adding a system

Cluster Explorer and Command Center enable you to add a system to the cluster. A system must have an entry in the LLTTab configuration file before it can be added to the cluster.

To add a system from Cluster Explorer

1. On the **Edit** menu, click **Add**, and click **System**.
   or
   Click **Add System** on the Cluster Explorer toolbar.
2. Enter the name of the system.
3. Click **Show Command** in the bottom left corner to view the command associated with the system. Click **Hide Command** to close the view of the command.
4. Click **OK**.

To add a system from Command Center

1. Click **Add System** in the Command Center toolbar.

or

In the Command Center configuration tree, expand **Commands** > **Configuration** > **Cluster Objects** > **Add System**.

2 Enter the name of the system.

3 Click **Apply**.

Deleting a system

To delete a system from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Cluster Objects** > **Delete System**.

2 Click the system.

3 Click **Apply**.

Freezing a system

Freeze a system to prevent service groups from coming online on the system.

To freeze a system from Cluster Explorer

1. Click the **Systems** tab of the configuration tree.
2. In the configuration tree, right-click the system, click **Freeze**, and click **Temporary** or **Persistent** from the menu. The persistent option maintains the frozen state after a reboot if the user saves this change to the configuration.

To freeze a system from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Operations** > **Availability** > **Freeze System**.
2. Click the system.
3. If necessary, select the **persistent** and **evacuate** check boxes. The evacuate option moves all service groups to a different system before the freeze operation takes place. The persistent option maintains the frozen state after a reboot if the user saves this change to the configuration.
4. Click **Apply**.

Unfreezing a system

Unfreeze a frozen system to enable service groups to come online on the system.

To unfreeze a system from Cluster Explorer

1. Click the **Systems** tab of the configuration tree.
2. In the configuration tree, right-click the system and click **Unfreeze**.

To unfreeze a system from Command Center

1. In the Command Center configuration tree, expand **Commands > Operations** > **Availability** > **Unfreeze System**.
2. Click the system.
3. Click **Apply**.

Administering clusters

Use the Java Console to specify the clusters you want to view from the console, and to modify the VCS configuration. The configuration details the parameters of the entire cluster. Use Cluster Explorer or Command Center to open, save, and “save and close” a configuration. VCS Simulator enables you to administer the configuration on the local system while VCS is offline.

Opening a cluster configuration

Use Cluster Explorer or Command Center to open or make changes to the VCS configuration.

To open a configuration from Cluster Explorer

On the File menu, click **Open Configuration**.

or

Click **Open Configuration** on the Cluster Explorer toolbar.

To open a configuration from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Configuration** > **Configuration File** > **Open Configuration**.
2. Click **Apply**.

Saving a cluster configuration

After updating the VCS configuration, use Cluster Explorer or Command Center to save the latest configuration to disk while maintaining the configuration state in read-write mode.

To save a configuration from Cluster Explorer

On the **File** menu, click **Save Configuration**.

or

Click **Save Configuration** on the Cluster Explorer toolbar.

To save a configuration from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Configuration** > **Configuration File** > **Save Configuration**.
2. Click **Apply**.

Saving and closing a cluster configuration

After updating the VCS configuration, use Cluster Explorer or Command Center to save the latest configuration to disk, and close or change the configuration state to read-only mode.

To save and close a configuration from Cluster Explorer

On the **File** menu, click **Close Configuratio**n.

or

Click **Save and Close Configuration** on the Cluster Explorer toolbar.

To save and close a configuration from Command Center

1. In the Command Center configuration tree, expand **Commands** > **Configuration** > **Configuration File** > **Close Configuration**.
2. Click **Apply**.

Executing commands

Use Command Center to execute commands on a cluster. Command Center enables you to run commands organized as “Configuration” and “Operation.”

To execute a command from Command Center

1. From Command Center, click the command from the command tree. If necessary, expand the tree to view the command.
2. In the corresponding command interface, click the VCS objects and appropriate options (if necessary).
3. Click **Apply**.

Editing attributes

Use the Java Console to edit attributes of VCS objects. By default, the Java Console displays key attributes and type specific attributes. To view all attributes associated with an object, click **Show all attributes**.

To edit an attribute from Cluster Explorer

1. From the Cluster Explorer configuration tree, click the object whose attributes you want to edit.
2. In the view panel, click the **Properties** tab. If the attribute does not appear in the Properties View, click **Show all attributes**. This opens the Attributes View.

3 In the Properties or Attributes View, click the icon in the **Edit** column of the **Key Attributes** or **Type Specific Attributes** table. In the Attributes View, click the icon in the **Edit** column of the attribute.

4 In the Edit Attribute dialog box, enter the changes to the attribute values.
To edit a scalar value:
Enter or click the value.
To edit a non-scalar value:
Use the **+** button to add an element. Use the **-** button to delete an element.
To change the attribute's scope:
Click the **Global** or **Per System** option.
To change the system for a local attribute:
Click the system from the menu.

5 Click **OK**.

To edit an attribute from Command Center

1 In the Command Center configuration tree, expand **Commands** > **Configuration** > **Attributes** > **Modify *vcs_object* Attributes**.

2 Click the VCS object from the menu.

3 In the attribute table, click the icon in the **Edit** column of the attribute.

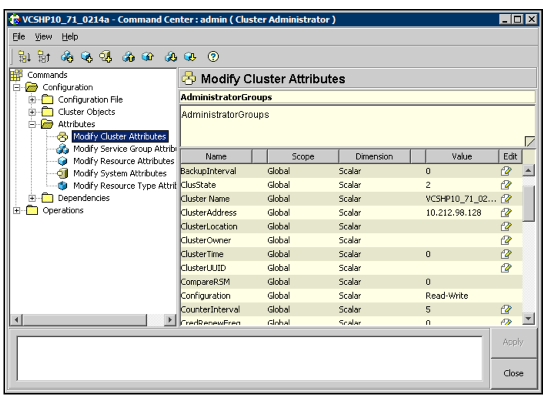

4 In the Edit Attribute dialog box, enter the changes to the attribute values.
To edit a scalar value:
Enter or click the value.
To edit a non-scalar value:
Use the **+** button to add an element. Use the **-** button to delete an element.
To change the attribute's scope:
Click the **Global** or **Per System** option.
To change the system for a local attribute:

Click the system from the menu.

5 Click **OK**.

Querying the cluster configuration

1 From Cluster Explorer, click **Query** on the **Tools** menu.
or
On the Cluster Explorer toolbar, click **Query**.

2 Enter the details of the query:

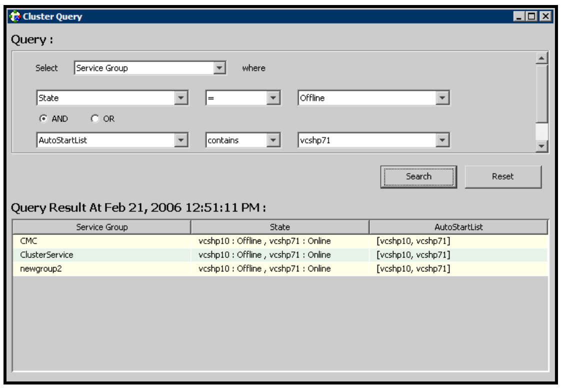

- Click the VCS object to search.
- Depending on the selected object, click the specific entity to search.
- Click the appropriate phrase or symbol between the search item and value.
- Click the appropriate value for the specified query. Certain queries allow the user to enter specific filter information:
  Click **System**, click **Online Group Count**, click **<**, and type the required value in the blank field.
  or
  Click **Resource**, click **[provide attribute name]** and type in the name of an attribute, click **=** or **contains**, and type the appropriate value of the attribute in the blank field. For example, click **Resource**, click **[provide attribute name]** and type in pathname, click **contains**, and type **c:\temp** in the blank field.
- To use additional queries, click **+** as many times as necessary to select the appropriate options. Click **-** to reduce the number of queries.
- Click **AND** or **OR** for each filter selection.
- Click **Search.** The results appear in tabular format at the bottom of the dialog box. To search a new item, click **Reset** to reset the dialog box to its original blank state.

Setting up VCS event notification using the Notifier wizard

The information presented here assumes that you need to create both the ClusterService group and the Notifier resource. If the ClusterService group exists but the Notifier resource is configured under another group, you can modify the attributes of the existing Notifier resource and system list for that group. If the ClusterService group is configured but the Notifier resource is not configured, the Notifier resource will be created and added to the ClusterService group.

To set up event notification using the Notifier wizard

1. From Cluster Explorer, click **Notifier Wizard...** on the **Tools** menu.
   or
   On the Cluster Explorer toolbar, click **Launch Notifier Resource Configuration Wizard**.
2. Click **Next**.
3. In the dialog box:

   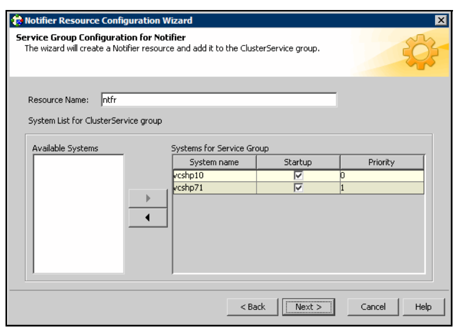

   - Enter the name of the resource. For example, "ntfr".
   - Click the target systems in the **Available Systems** box.
   - Click the right arrow to move the systems to the **Systems for Service Group** table. To remove a system from the table, click the system and click the left arrow.
   - Select the **Startup** check box to add the systems to the service groups AutoStartList attribute. This enables the service group to automatically come online on a system every time HAD is started.
   - The priority number (starting with 0) is assigned to indicate the order of systems on which the service group will start in case of a failover. If

necessary, double-click the entry in the **Priority** column to enter a new value.

- Click **Next**.

4 Choose the mode of notification which needs to be configured. Select the check boxes to configure SNMP and/or SMTP (if applicable).

5 In the SNMP Configuration dialog box (if applicable):

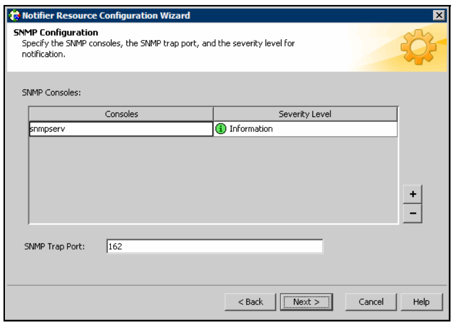

- Click **+** to create the appropriate number of fields for the SNMP consoles and severity levels. Click **-** to remove a field.
- Enter the console and click the severity level from the menu. For example, "snmpserv" and "Information".
- Enter the SNMP trap port. For example, "162" is the default value.
- Click **Next**.

6 In the SMTP Configuration dialog box (if applicable):

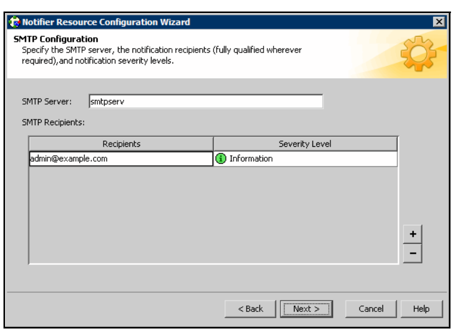

- Enter the name of the SMTP server.
- Click + to create the appropriate number of fields for recipients of the notification and severity levels. Click - to remove a field.

- Enter the recipient and click the severity level in the drop-down list box. For example, “admin@example.com” and “Information”.
- Click **Next**.

7 In the NIC Resource Configuration dialog box:

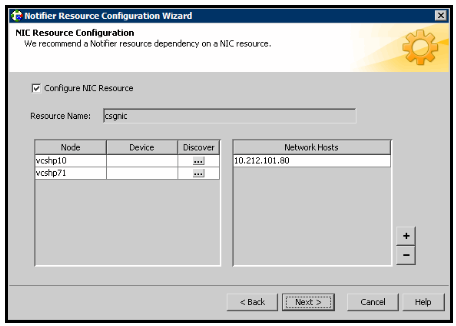

- Click **Configure NIC Resource** (recommended by Symantec) and proceed to the next step. Otherwise, click **Next**.
- If necessary, enter the name of the resource.
- Click the icon (**...**) in the **Discover** column of the table to find the MACAddress for each system.
- Click **OK** on the Discover dialog box.
- Click **Next**.

8 Click the **Bring the Notifier Resource Online** check box, if desired.

9 Click **Next**.

10 Click **Finish**.

Administering logs

The Java Console enables you to customize the log display of messages generated by the engine. In the Logs dialog box, you can set filter criteria to search and view messages, and monitor and resolve alert messages.

To browse the logs for detailed views of each log message, double-click the event's description. Use the arrows in the **VCS Log details** pop-up window to navigate backward and forward through the message list.

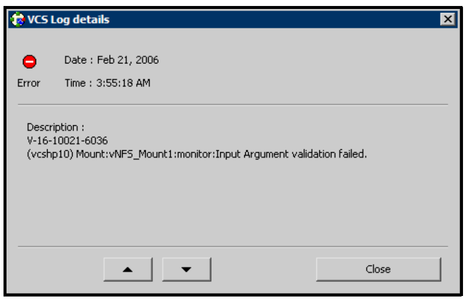

Customizing the log display

From the Logs dialog box, use the **Edit Filters** feature to customize the display of log messages.

To customize the display for VCS logs

1. In the **VCS Logs** tab, click **Edit Filters**.
2. Enter the filter criteria:
   - Click the types of logs to appear on the message display.

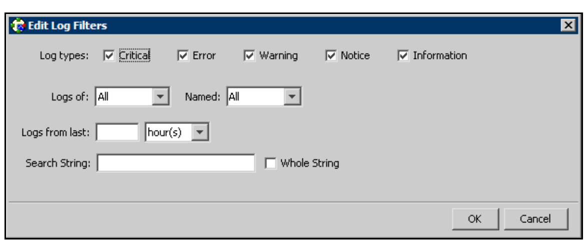

   - From the **Logs of** list, select the category of log messages to display.
   - From the **Named** menu, select the name of the selected object or component. To view all the messages for the selected category, click **All**.

- In the **Logs from last** field, enter the numerical value and select the time unit.
- To search log messages, enter the search string. Select the **Whole String** check box, if required.
- Click **OK**.

To customize the display for agent logs

1 In the **Agent Logs** tab, enter the filter criteria:

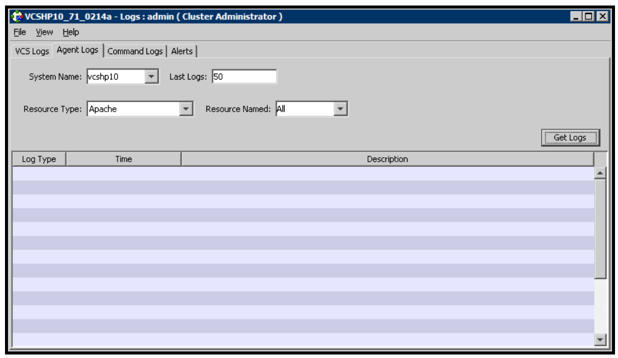

- Click the name of the system.
- Enter the number of logs to view.
- Click the resource type.
- Click the name of the resource. To view messages for all resources, click **All**.
- Click **Get Logs**.

Resetting the log display

Use the **Reset Filters** feature to set the default settings for the log view. For example, if you customized the log view to only show critical and error messages using the **Edit Filters** feature, the **Reset Filters** feature will set the view to show all log messages.

To reset the default settings for the log display

In the **VCS Logs** tab, click **Reset Filters**.

Monitoring alerts

The Java Console sends automatic alerts that require administrative action and are displayed on the **Alerts** tab of the Logs dialog box. Use this tab to take action on the alert or delete the alert.

To take action on an alert

1 In the **Alert** tab or dialog box, click the alert to take action on.

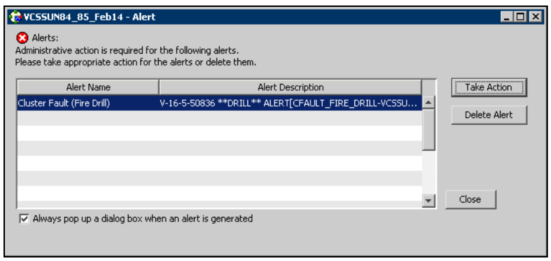

2 Click **Take Action**.

3 Enter the required information to resolve the alert.
If the alert warns that a local group cannot fail over to any system in the local cluster, you cannot take any action.
If the alert warns that a global group cannot fail over, the action involves bringing the group online on another system in the global cluster environment.
If the alert warns that a global cluster is faulted, the action involves declaring the cluster as a disaster, disconnect, or outage, and determining the service groups to fail over to another cluster.

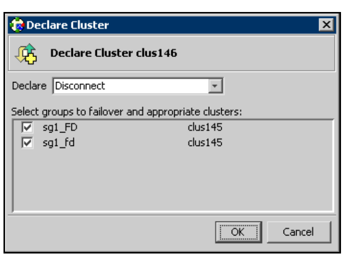

4 Click **OK**.

To delete an alert

1. In the **Alert** tab or dialog box, click the alert to delete.
2. Click **Delete Alert**.
3. Provide the details for this operation:

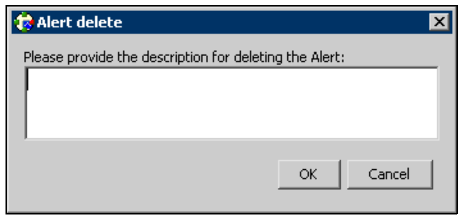

- Enter the reason for deleting the alert.
- Click **OK**.

Administering VCS Simulator

VCS Simulator enables you to view state transitions, experiment with configuration parameters, and predict how service groups will behave during cluster or system faults. Use this tool to create and save configurations in an OFFLINE state.

Through the Java Console, VCS Simulator enables you to configure a simulated cluster panel, bring a system in an unknown state into an online state, simulate power loss for running systems, simulate resource faults, and save the configuration while VCS is offline.

For global clusters, you can simulate the process of generating and clearing cluster faults.

You can run multiple simulated clusters on a system by using different port numbers for each cluster. The Java Console provides the same views and features that are available for online configurations

See "Predicting VCS behavior using VCS Simulator" on page 405.

Chapter 8

Administering the cluster from the command line

- About administering VCS from the command line
- Starting VCS
- Stopping VCS
- Managing VCS configuration files
- Managing VCS users from the command line
- Querying VCS
- Administering service groups
- Administering agents
- Administering resource types
- Administering resources
- Administering systems
- Administering clusters
- Administering simulated clusters from the command line

About administering VCS from the command line

This chapter describes commonly used VCS commands. For more information about specific commands or their options, see their usage information or the man pages associated with the commands.

Most commands can be entered from any system in the cluster when VCS is running. The command to start VCS is typically invoked at system startup.

How VCS identifies the local system

VCS uses the system's node name. To view the system's node name from the command line, type:

```
C:\> hostname
```

To view the system's node name from the desktop

1 Right-click Network Neighborhood to display the pop-up menu.

2 Click **Properties**. The name of the system is listed in the Identification tab.

About specifying values preceded by a dash (-)

When specifying values in a command-line syntax, you must prefix values beginning with a dash (-) with a percentage sign (%). If a value begins with a percentage sign, you must prefix it with another percentage sign. (The initial percentage sign is stripped by HAD and does not appear in the configuration file.)

About the -modify option

Most configuration changes are made using the `-modify` options of the commands `haclus`, `hagrp`, `hares`, `hasys`, and `hatype`. Specifically, the `-modify` option of these commands changes the attribute values stored in the VCS configuration file. By default, all attributes are global, meaning that the value of the attribute is the same for all systems.

Note: VCS must be in read/write mode before you can change the configuration. For instructions, see "Setting the configuration to read/w rite" on page 242.

Encrypting passwords

Use the `vcsencrypt` utility to encrypt passwords when editing the VCS configuration file main.cf to add VCS users or when configuring agents that require user passwords.

Note: Do not use the `vcsencrypt` utility when entering passwords from a configuration wizard or from the Java and Web consoles.

To encrypt a password

1. Run the utility from the command line.
   To encrypt a password for an agent configuration:
   `vcsencrypt -agent`
   To encrypt a VCS user password:
   `vcsencrypt -vcs`
2. The utility prompts you to enter the password twice. Enter the password and press Return.
   ```
   Enter New Password:
   Enter Again:
   ```
3. The utility encrypts the password and displays the encrypted password. Use the displayed password to edit the VCS configuration file main.cf.

Starting VCS

When VCS is started, it checks the state of its local configuration file and registers with GAB for cluster membership. If the local configuration is valid, and if no other system is running VCS, it builds its state from the local configuration file and enters the RUNNING state.

If the configuration on all nodes is invalid, the VCS engine waits for manual intervention, or for VCS to be started on a system that has a valid configuration.

See "Cluster and system states" on page 693.

To start VCS

- Run the following command:
  `hastart`

To start VCS when all systems are in the ADMIN_WAIT state

- Run the following command from any system in the cluster to force VCS to use the configuration file from the system specified by the variable *system*:
  `hasys -force` ***`system`***

To start VCS on a single node

- Type the following command to start an instance of VCS that does not require the GAB and LLT packages. Do not use this command on a multisystem cluster.

hastart -onenode

To start VCS as a time-sharing process

- Run the following command:

  hastart -ts

To start CommandServer

- Run the following command:

  net start cmdserver

Stopping VCS

The `hastop` command stops HAD and related processes. You can customize the behavior of the hastop command by configuring the EngineShutdown attribute for the cluster.

See “Controlling the hastop behavior using the EngineShutdown attribute” on page 239.

The hastop command includes the following options:

```
hastop -all [-force]
hastop [-help]
hastop -local [-force | -evacuate | -noautodisable]
hastop -sys system ... [-force | -evacuate | -noautodisable]
```

Option	Description
`-all`	Stops HAD on all systems in the cluster and takes all service groups offline.
`-help`	Displays command usage.
`-local`	Stops HAD on the system on which you typed the command
`- force`	Allows HAD to be stopped without taking service groups offline on the system. The value of the EngineShutdown attribute does not influence the behavior of the -force option.
`-evacuate`	When combined with `-local` or `-sys`, migrates the system’s active service groups to another system in the cluster, before the system is stopped.
`-noautodisable`	Ensures that service groups that can run on the node where the `hastop` command was issued are not autodisabled. This option can be used with `evacuate` but not with `-force`.
`-sys`	Stops HAD on the specified system.

Stopping VCS without -force option

When VCS is stopped on a system without using the `-force` option, it enters the LEAVING state, and waits for all groups to go offline on the system. Use the output of the command `hasys -display` *system* to verify that the values of the SysState and the OnGrpCnt attributes are non-zero. VCS continues to wait for the service groups to go offline before it shuts down.

See "Troubleshooting resources" on page 625.

Stopping VCS with options other than -force

When VCS is stopped by options other than `-force` on a system with online service groups, the groups running on the system are taken offline and remain offline. This is indicated by VCS setting the attribute IntentOnline to 0. Using the option `-force` enables service groups to continue running while HAD is brought down and restarted (IntentOnline remains unchanged).

Controlling the hastop behavior using the EngineShutdown attribute

Use the EngineShutdown attribute to define VCS behavior when a user runs the hastop command.

Note: VCS does not consider this attribute when the hastop is issued with the following options: -force or -local -evacuate -noautodisable.

Configure one of the following values for the attribute depending on the desired functionality for the hastop command:

EngineShutdown Value	Description
Enable	Process all hastop commands. This is the default behavior.
Disable	Reject all hastop commands.
DisableClusStop	Do not process the hastop -all command; process all other hastop commands.
PromptClusStop	Prompt for user confirmation before running the hastop -all command; process all other hastop commands.
PromptLocal	Prompt for user confirmation before running the hastop -local command; process all other hastop commands.
PromptAlways	Prompt for user confirmation before running any hastop command.

Additional considerations for stopping VCS

- If using the command `reboot`, behavior is controlled by the ShutdownTimeOut parameter. After HAD exits, if GAB exits within the time designated in the ShutdownTimeout attribute, the remaining systems recognize this as a reboot and fail over service groups from the departed system. For systems running several applications, consider increasing the value in the ShutdownTimeout attribute.
- Stopping VCS on a system autodisables each service group that includes the system in their SystemList attribute. (This does not apply to systems that are powered off.)
- If you use the `-evacuate` option, evacuation occurs before VCS is brought down.

Managing VCS configuration files

This section describes how to verify, back up, and restore VCS configuration files.

See "About the main.cf file" on page 49.
See "The types.cf file" on page 51.

About the hacf utility

The hacf utility translates the VCS configuration language into a syntax that can be read by the VCS engine. Specifically, hacf translates the contents of the main configuration file, main.cf, into commands for the VCS server.

The hacf utility verifies the configuration before loading it into VCS. The configuration is not loaded under the following conditions:

- If main.cf or include files are missing.
- If syntax errors appear in the .cf files.
- If the configuration file is invalid.

See "Setting the Configuration to Read/Write on page 67.

About multiple versions of .cf files

When hacf creates a .cf file, it does *not* overwrite existing .cf files. A copy of the file remains in the directory, and its name includes a suffix of the date and time it was created, such as main.cf.03Dec2001.175904. In addition, the previous version of any .cf file is saved with the suffix .previous; for example, main.cf.previous.

Verifying a configuration

Use hacf to verify (check syntax of) the main.cf and the type definition file, types.cf. VCS does not execute if hacf detects errors in the configuration.

To verify a configuration

◆ Run the following command:

```
hacf -verify config_directory
```

The variable *config_directory* refers to directories containing a main.cf file and any .cf files included in main.cf.

No error message and a return value of zero indicates that the syntax is legal.

Scheduling automatic backups for VCS configuration files

Configure the BackupInterval attribute to instruct VCS to create a back up of the configuration periodically. VCS backs up the main.cf and types.cf files as main.cf.autobackup and types.cf.autobackup respectively.

To start periodic backups of VCS configuration files

◆ Set the cluster-level attribute BackupInterval to a non-zero value.
  For example, to back up the configuration every 5 minutes, set BackupInterval to 5.

Saving a configuration

When you save a configuration, VCS renames the file main.cf.autobackup to main.cf. VCS also save your running configuration to the file main.cf.autobackup.

If have not configured the BackupInterval attribute, VCS saves the running configuration.

See "Scheduling automatic backups for VCS configuration files" on page 242.

To save a configuration

◆ Run the following command

```
haconf -dump -makero
```

  The option -makero sets the configuration to read-only.

Setting the configuration to read/write

To set the mode to read/write

◆ Type the following command:

```
haconf -makerw
```

Taking snapshots of VCS configuration files

Use the hasnap command to take snapshots of VCS configuration files on each node in the cluster. You can also restore the configuration from a snapshot.

The command includes the following options; each option is described in detail in the following sections:

`hasnap -backup`	Backs up files in a snapshot format.
`hasnap -restore`	Restores a previously created snapshot.
`hasnap -display`	Displays details of previously created snapshots.
`hasnap -sdiff`	Displays files that were changed on the local system after a specific snapshot was created.
`hasnap -fdiff`	Displays the differences between a file in the cluster and its copy stored in a snapshot.
`hasnap -export`	Exports a snapshot from the local, predefined directory to the specified file.
`hasnap -include`	Configures the list of files or directories to be included in new snapshots, in addition to those included automatically by the `-backup` command.
`hasnap -exclude`	Configures the list of files or directories to be excluded from new snapshots when backing up the configuration using the `-backup` command.
`hasnap -delete`	Deletes snapshots from the predefined local directory on each node.

Note: With the exception of the `-include`, `-exclude`, and the `-delete` options, all options can be combined with the `-f` option. This option indicates that all files be backed up to or restored from the specified single file instead of a local, predefined directory on each node. This option is useful when you want to store the configuration data to an alternate location that is periodically backed up using backup software like Veritas Net Backup.

Backing up configuration files

The `hasnap -backup` command backs up files in a snapshot format. A snapshot is a collection of VCS configuration files backed up at a particular point in time, typically before making changes to the existing configuration. A snapshot also contains information such as the snapshot name, description, creation time, and file permissions.

The command backs up a predefined list of VCS configuration files as well as a user-defined list. The predefined list includes all the *.cf files, custom agents, LLT and GAB configuration files, triggers, custom heartbeats, and action scripts. See the `-include` and `-exclude` commands to construct a user-defined list.

To back up VCS configuration files

- Run the following command
  hasnap -backup [-f *filename*] [-n] [-m *description*]

  Use the `-n` option to run the command in the non-interactive mode.
  Use the `-m` option to specifies a description of the snapshot.

Examples

The following command creates a backup of the configuration in the non-interactive mode and adds Test Backup as the backup description.

```
hasnap -backup -n -m "Test Backup"
```

Restoring VCS configuration files

The `hasnap -restore` command restores configuration files from a previously created snapshot.

To restore VCS configuration files

- Run the following command:
  hasnap -restore [-f *filename*] [-n] [-s *snapid*]

`-n`	Run the command in the non-interactive mode.
`-s`	option to specifies the ID of the snapshot to be restored.

  If you do not specify a snapshot ID, the command lists the snapshots that are available for restoration.

Examples

The following command restores the snapshot vcs-20030101-22232 in the non-interactive mode.

```
hasnap -restore -n -s vcs-20030101-22232
```

Viewing snapshots of configuration files

Use the `hasnap -display` command to view details of previously created snapshots.

To view snapshots of configuration files

hasnap -display [-f *filename*] [-list|-s *snapid*] [-m] [-l] [-t]

`-list`	Displays the list of snapshots in the repository.
`-s`	Specifies the snapshot ID.
`-m`	Displays snapshot description.
`-l`	Lists files in the snapshot
`-t`	Displays the snapshot timestamp

If no options are specified, the command displays all information about the latest snapshot.

Examples

The following command lists all snapshots.

hasnap -display -list

The following command displays the description and the time of creation of the specified snapshot.

hasnap -display -s vcs-20030101-2232 -m -t

Viewing files changed after a snapshot

Use the `hasnap -sdiff` command to display files that were changed on the local system after a specific snapshot was created.

To view files that changed after a snapshot

◆ Run the following command:

hasnap -sdiff [-f *filename*] [-s *snapid*] [-sys *hostname*]

`-s`	Identifies the snapshot ID of the comparison snapshot.
`-sys`	Indicates the host on which the snapshot is to be compared.

If you do not specify any options, the command uses the latest snapshot to compare the files on each node in the cluster.

Examples

The following command displays the differences between the current configuration and the snapshot vcs-20030101-22232.

```
hasnap -sdiff -s vcs-20030101-22232
```

Comparing a file with its snapshot copy

Use the `hasnap -fdiff` to displays differences between a file on the cluster and its copy stored in a previously created snapshot.

To compare a file with its snapshot copy

◆ Run the following command:
hasnap -fdiff [-f *filename*] [-s *snapid*] [-sys *hostname*] file

`-s`	Specifies the ID of the snapshot.
`-sys`	Specifies the host on which the snapshot is to be compared.
`-file`	The file to compare.

If you do not specify any options, the command uses the latest snapshot to compare the file on each node in the cluster.

Exporting snapshots

Use the `hasnap -export` command to export a snapshot from the local, predefined directory on each node in the cluster to a specified file. This option is useful when you want to store a previously created snapshot to an alternate location that is periodically backed up using backup software like Veritas NetBackup.

To export a snapshot

◆ Run the following command:
hasnap -export -f *filename* [-s *snapid*]

`-s`	Specifies the ID of the snapshot.
`-f`	Specifies the file.

If you do not specify a snapshot ID, the command exports the latest snapshot to the specified file.

Adding and removing files for snapshots

Use the `hasnap -include` command to configures the list of files or directories to be included in new snapshots. This list is in addition to the files included by the `-backup` command.

See "Backing up configuration files" on page 244.

To add or remove files for a snapshots

- Run the following command:

  ```
  hasnap -include -add|-del|-list [-sys hostname]
      files|directories
  ```

`-add`	Adds the specified files or directories to the include file list.
`-del`	Removes the specified files or directories from the include file list.
`files/directories`	Files or directories to be added or removed.

Examples

The following command displays the list of files or directories to be included in new snapshots on each node of the cluster.

```
hasnap -include -list
```

Excluding files from snapshots

Use the `hasnap -exclude` command to configure the list of files or directories that should not be included in new snapshots.

To exclude files from snapshots

- Run the following command:

  ```
  hasnap -exclude -add|-del|-list [-sys hostname]
      files|directories
  ```

`-add`	Adds the specified files or directories to the exclude file list.
`-del`	Removes the specified files or directories from the exclude file list.
`files/directories`	Files or directories to be added or removed.

Examples

The following command displays the exclude file list on each node in the cluster.

```
hasnap -exclude -list
```

Deleting snapshots

Use the `hasnap -delete` command to delete snapshots from the predefined local directory on each node.

To delete a snapshot

- Run the following command:

  `hasnap -delete [-s` ***`snapid`*****`]`**

`-s`	Specifies the ID of the snapshot to be deleted.

  If you do not specify the snapshot ID, the command lists the snapshots that can be deleted.

Example

The following command deletes snapshot vcs-20030101-22232 from the cluster.

```
hasnap -delete -s vcs-20030101-22232
```

Managing VCS users from the command line

You can add, modify, and delete users on any system in the cluster, provided you have the privileges to do so.

If VCS is running in secure mode, 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.

The commands to add, modify, and delete a user must be executed only as root or administrator and only if the VCS configuration is in read/write mode.

See "Setting the configuration to read/w rite" on page 242.

Note: You must add users to the VCS configuration to monitor and administer VCS from the graphical user interface Cluster Manager.

Adding a user

Users in the category Cluster Guest cannot add users.

To add a user

1 Set the configuration to read/write mode:

```
haconf -makerw
```

2 Add the user:

```
hauser -add user [-priv <Administrator|Operator> [-group
      service_groups]]
```

3 Enter a password when prompted.

4 Reset the configuration to read-only:

```
haconf -dump -makero
```

To add a user with cluster administrator access

◆ Type the following command:

```
hauser -add user -priv Administrator
```

To add a user with cluster operator access

◆ Type the following command:

```
hauser -add user -priv Operator
```

To add a user with group administrator access

◆ Type the following command:

```
hauser -add user -priv Administrator -group service_groups
```

To add a user with group operator access

◆ Type the following command:

```
hauser -add user -priv Operator -group service_groups
```

Assigning and removing user privileges

To assign privileges to an administrator or operator

◆ Type the following command:

```
hauser -addpriv user Adminstrator|Operator
       [-group service_groups]
```

To remove privileges from an administrator or operator

◆ Type the following command:

```
hauser -delpriv user Adminstrator|Operator
               [-group service_groups]
```

To assign privileges to an OS user group

◆ Type the following command:

```
hauser -addpriv usergroup AdminstratorGroup|OperatorGroup
       [-group service_groups]
```

To remove privileges from an OS user group

◆ Type the following command:

```
hauser -delpriv usergroup AdminstratorGroup|OperatorGroup
   [-group service_groups]
```

Modifying a user

Users in the category Cluster Guest cannot modify users.

To modify a user

1 Set the configuration to read/write mode:

```
haconf -makerw
```

2 Modify the user:

```
hauser -update user
```

3 Enter a new password when prompted.

4 Reset the configuration to read-only:

```
haconf -dump -makero
```

Deleting a user

You can delete a user from the VCS configuration.

To delete a user

1. Set the configuration to read/write mode:

   ```
   haconf -makerw
   ```

2. For users with Administrator and Operator access, remove their privileges:

   ```
   hauser -delpriv user Adminstrator|Operator [-group
       service_groups]
   ```

3. Delete the user from the list of registered users:

   ```
   hauser -delete user
   ```

4. Reset the configuration to read-only:

   ```
   haconf -dump -makero
   ```

Displaying a user

Display a list of users and their privileges.

To display a list of users

- Type the following command:

  ```
  hauser -list
  ```

To display the privileges of all users

- Type the following command:

  ```
  hauser -display
  ```

To display the privileges of a specific user

- Type the following command:

  ```
  hauser -display user
  ```

Querying VCS

VCS enables you to query various cluster objects, including resources, service groups, systems, resource types, agents, and clusters. You may enter query commands from any system in the cluster. Commands to display information on the VCS configuration or system states can be executed by all users: you do not need root privileges.

Querying service groups

To display the state of a service group on a system

- Type the following command:
  hagrp -state [*service_group*] [-sys *system*]

To display the resources for a service group

- Type the following command:
  hagrp -resources *service_group*

To display a list of a service group's dependencies

- Type the following command:
  hagrp -dep [*service_group*]

To display a service group on a system

- Type the following command:
  hagrp -display [*service_group*] [-sys *system*]
  If *service_group* is not specified, information regarding all service groups is displayed.

To display attributes of a system

- Type the following command:
  **hagrp -display [*service_group*] [-attribute *attribute*]
  [-sys *system*]**

  Note that system names are case-sensitive.

Querying resources

To display a resource's dependencies

- Type the following command:

  hares -dep [*resource*]

To display information about a resource

- Type the following command:

  hares -display [*resource*]

  If *resource* is not specified, information regarding all resources is displayed.

To confirm an attribute's values are the same on all systems

- Type the following command:

  hares -global *resource* *attribute* *value* ... | *key*... | {*key* *value*}...

To display resources of a service group

- Type the following command:

  hares -display -group *service_group*

To display resources of a resource type

- Type the following command:

  hares -display -type *resource_type*

To display attributes of a system

- Type the following command:

  hares -display -sys *system*

Querying resource types

To display all resource types

- Type the following command:

  hatype -list

To display resources of a particular resource type

- Type the following command:

  hatype -resources *resource_type*

To display information about a resource type

- Type the following command:
- Type the following command:

  hatype -display *resource_type*

  If *resource_type* is not specified, information regarding all types is displayed.

Querying agents

To display the run-time status of an agent'

- Type the following command:

  haagent -display [*agent*]

  If *agent* is not specified, information regarding all agents is displayed.

Run-Time Status	Definition
Faults	Indicates the number of agent faults and the time the faults began.
Messages	Displays various messages regarding agent status.
Running	Indicates the agent is operating.
Started	Indicates the file is executed by the VCS engine (HAD).

Querying systems

To display a list of systems in the cluster

- Type the following command:
  hasys -list

To display information about each system

- Type the following command:
  hasys -display [*system*]

Querying clusters

To display the value of a specific cluster attribute

- Type the following command:
  haclus -value *attribute*

To display information about the cluster

- Type the following command:
  haclus -display

Querying status

To display the status of all service groups in the cluster, including resources

- Type the following command:
  hastatus

To display the status of a particular service group, including its resources

- Type the following command:
  hastatus [-sound] -group *service_group* [-group *service_group*]...
  If you do not specify a service group, the status of all service groups is displayed. The `-sound` option enables a bell to ring each time a resource faults.

To display the status of service groups and resources on specific systems

- Type the following command:
  hastatus [-sound] -sys *system_name* [-sys *system_name*]...

To display the status of specific resources

◆ Type the following command:

hastatus [-sound] -resource *resource_name* [-resource *resource_name*]...

To display the status of cluster faults, including faulted service groups, resources, systems, links, and agents

◆ Type the following command:

hastatus -summary

Note: Unless executed with the `-summary` options, the `hastatus` command continues to produce output of online state transitions until you interrupt it with the command CTRL+C.

Querying log data files (LDFs)

Log data files (LDFs) contain data regarding messages written to a corresponding English language file. Typically, for each English file there is a corresponding LDF.

To display the hamsg usage list

◆ Type the following command:

hamsg -help

To display the list of LDFs available on the current system

◆ Type the following command:

hamsg -list

To display general LDF data

◆ Type the following command:

hamsg -info [-path *path_name*] *LDF*

The option `-path` specifies where `hamsg` looks for the specified LDF. If not specified, `hamsg` looks for files in the default directory:

Program Files\VERITAS\Cluster Server\ldf

To display specific LDF data

◆ Type the following command:

hamsg [-any] [-sev *C*|*E*|*W*|*N*|*I*] [-otype *VCS*|*RES*|*GRP*|*SYS*|*AGT*] [-oname *object_name*] [-msgid *message_ID*] [-path *path_name*] [-lang *language*] *LDF*

`-any`	Specifies `hamsg` return messages matching any of the specified query options.
`-sev`	Specifies `hamsg` return messages matching the specified message severity Critical, Error, Warning, Notice, or Information.
`-otype`	Specifies `hamsg` return messages matching the specified object type ■ VCS = general VCS messages ■ RES = resource ■ GRP = service group ■ SYS = system ■ AGT = agent
`-oname`	Specifies `hamsg` return messages matching the specified object name.
`-msgid`	Specifies `hamsg` return messages matching the specified message ID.
`-path`	Specifies where `hamsg` looks for the specified LDF. If not specified, `hamsg` looks for files in the default directory /var/VRTSvcs/ldf.
`-lang`	Specifies the language in which to display messages. For example, the value en specifies English and "ja" specifies Japanese.

Using conditional statements to query VCS objects

Some query commands include an option for conditional statements. Conditional statements take three forms:

`Attribute=Value` (the attribute equals the value)

`Attribute!=Value` (the attribute does not equal the value)

`Attribute=~Value` (the value is the prefix of the attribute, for example a query for the state of a resource = ~FAULTED returns all resources whose state begins with FAULTED.)

Multiple conditional statements can be used and imply AND logic.

You can only query attribute-value pairs displayed in the output of the command `hagrp -display`.

See "Querying service groups" on page 252.

To display the list of service groups whose values match a conditional statement

- Type the following command:

  hagrp -list [*conditional_statement*]

  If no conditional statement is specified, all service groups in the cluster are listed.

To display a list of resources whose values match a conditional statement

- Type the following command:

  hares -list [*conditional_statement*]

  If no conditional statement is specified, all resources in the cluster are listed.

To display a list of agents whose values match a conditional statement

- Type the following command:

  haagent -list [*conditional_statement*]

  If no conditional statement is specified, all agents in the cluster are listed.

Administering service groups

This section describes how to add, delete, and modify service groups. It also describes how to perform service group operations from the command line.

Adding and deleting service groups

To add a service group to your cluster

hagrp -add *service_group*

The variable *service_group* must be unique among all service groups defined in the cluster.

This command initializes a service group that is ready to contain various resources. To employ the group properly, you must populate its SystemList attribute to define the systems on which the group may be brought online and taken offline. (A system list is an association of names and integers that represent priority values.)

To delete a service group

- Type the following command:

  hagrp -delete *service_group*

  Note that you cannot delete a service group until all of its resources are deleted.

Modifying service group attributes

To modify a service group attribute

- Type the following command:

  hagrp -modify *service_group attribute value* [-sys *system*]

  The variable *value* represents:

  system_name1 priority system_name2 priority2

  If the attribute being modified has local scope, you must specify the system on which to modify the attribute, except when modifying the attribute on the system from which you run the command.

  For example, to populate the system list of service group groupx with Systems A and B, type:

  ```
  hagrp -modify groupx SystemList -add SystemA 1 SystemB 2
  ```

  Similarly, to populate the AutoStartList attribute of a service group, type:

  ```
  hagrp -modify groupx AutoStartList SystemA SystemB
  ```

You may also define a service group as parallel. To set the Parallel attribute to 1, type the following command. (Note that the default for this attribute is 0, which designates the service group as a failover group.):

```
hagrp -modify groupx Parallel 1
```

This attribute cannot be modified if resources have already been added to the service group.

You can modify the attributes SystemList, AutoStartList, and Parallel only by using the command `hagrp -modify`. You cannot modify attributes created by the system, such as the state of the service group.

About modifying the SystemList attribute

When using the `hagrp -modify` command ot change a service group's existing system list, you can use the options `-modify`, `-add`, `-update`, `-delete`, or `-delete -keys`.

For example, suppose you originally defined the SystemList of service group groupx as SystemA and SystemB. Then after the cluster was brought up you added a new system to the list:

```
hagrp -modify groupx SystemList -add SystemC 3
```

You must take the service group offline on the system being modified.

When you add a system to a service group's system list, the system must have been previously added to the cluster. When using the command line, you can use the `hasys -add` command.

When you delete a system from a service group's system list, the service group must not be online on the system to be deleted.

If you attempt to change a service group's existing system list using `hagrp -modify` without other options (such as `-add` or `-update`) the command fails.

Bringing service groups online

To bring a service group online

- Type the following command:

  hagrp -online *service_group* -sys *system*

To start a service group on a system and bring online only the resources already online on another system

- Type the following command:

  hagrp -online *service_group* -sys *system* -checkpartial
  other_system

  If the service group does not have resources online on the other system, the service group is brought online on the original system and the `checkpartial` option is ignored.

  Note that the `checkpartial` option is used by the Preonline trigger during failover. When a service group configured with Preonline =1 fails over to another system (system 2), the only resources brought online on system 2 are those that were previously online on system 1 prior to failover.

Taking service groups offline

To take a service group offline

- Type the following command:

  hagrp -offline *service_group* -sys *system*

To take a service group offline only if all resources are probed on the system

- Type the following command:

  hagrp -offline [-ifprobed] *service_group* -sys *system*

Switching service groups

The process of switching a service group involves taking it offline on its current system and bringing it online on another system

To switch a service group from one system to another

- Type the following command:

  hagrp -switch *service_group* -to *system*

  A service group can be switched only if it is fully or partially online. The `-switch` option is not supported for switching parallel service groups and for switching hybrid service groups across system zones.

Freezing and unfreezing service groups

Freeze a service group to prevent it from failing over to another system. This freezing process stops all online and offline procedures on the service group.

Unfreeze a frozen service group to perform online or offline operations on the service group.

To freeze a service group (disable online, offline, and failover operations)

- Type the following command:

  hagrp -freeze *service_group* [-persistent]

  The option `-persistent` enables the freeze to be remembered when the cluster is rebooted.

To unfreeze a service group (reenable online, offline, and failover operations)

- Type the following command:

  hagrp -unfreeze *service_group* [-persistent]

Enabling and disabling service groups

Enable a service group before bringing it online. A service group that was manually disabled during a maintenance procedure on a system may need to be brought online after the procedure is completed.

Disable a service group to prevent it from coming online. This process temporarily stops VCS from monitoring a service group on a system undergoing maintenance operations

To enable a service group

- Type the following command:

  hagrp -enable *service_group* [-sys *system*]

  A group can be brought online only if it is enabled.

To disable a service group

- Type the following command:

  hagrp -disable *service_group* [-sys *system*]

  A group cannot be brought online or switched if it is disabled.

To enable all resources in a service group

- Type the following command:

  hagrp -enableresources *service_group*

To disable all resources in a service group

- Type the following command:

  hagrp -disableresources *service_group*

  Agents do not monitor group resources if resources are disabled.

Clearing faulted resources in a service group

Clear a resource to remove a fault and make the resource available to go online.

To clear faulted, non-persistent resources in a service group

- Type the following command:

  hagrp -clear *service_group* [-sys *system*]

  Clearing a resource initiates the online process previously blocked while waiting for the resource to become clear.

  - If *system* is specified, all faulted, non-persistent resources are cleared from that system only.
  - If *system* is not specified, the service group is cleared on all systems in the group's SystemList in which at least one non-persistent resource has faulted.

To clear resources in ADMIN_WAIT state in a service group

- Type the following command:

  hagrp -clearadminwait [-fault] *service_group* -sys *system*

  See "Changing agent file paths and binaries" on page 471.

Linking and unlinking service groups

Link service groups to create a dependency between them.

See "About service group dependencies" on page 492.

To link service groups

- Type the following command

  hagrp -link *parent_group child_group gd_category* *gd_location gd_type*

parent_group	Name of the parent group
child_group	Name of the child group
gd_category	category of group dependency (online/offline).
gd_location	the scope of dependency (local/global/remote).

gd_type type of group dependency (soft/firm/hard). Default is firm.

To unlink service groups

- Type the following command:

  hagrp -unlink *parent_group child_group*

Administering agents

Under normal conditions, VCS agents are started and stopped automatically.

To start an agent

- Run the following command:

  haagent -start *agent* -sys *system*

To stop an agent

- Run the following command:

  haagent -stop *agent* [-force] -sys *system*

The -force option stops the agent even if the resources for the agent are online. Use the -force option when you want to upgrade an agent without taking its resources offline.

Administering resources

Adding resources

Add resource to a service group or remove resources from a service group.

To add a resource

- Type the following command:

  hares -add *resource resource_type service_group*

The resource name must be unique throughout the cluster. The resource type must be defined in the configuration language. The resource belongs to the group *service_group*.

About adding resources

When you add a resource, all non-static attributes of the resource's type, plus their default values, are copied to the new resource.

Three attributes are also created by the system and added to the resource:

- Critical (default = 1). If the resource or any of its children faults while online, the entire service group is marked faulted and failover occurs.
- AutoStart (default = 1). If the resource is set to AutoStart, it is brought online in response to a service group command. All resources designated as AutoStart=1 must be online for the service group to be considered online. (This attribute is unrelated to AutoStart attributes for service groups.)
- Enabled. If the resource is set to Enabled, the agent for the resource's type manages the resource. The default is 1 for resources defined in the configuration file main.cf, 0 for resources added on the command line.

Note: Adding resources on the command line requires several steps, and the agent must be prevented from managing the resource until the steps are completed. For resources defined in the configuration file, the steps are completed before the agent is started.

Deleting resources

Delete resources from a service group.

To delete a resource

- Type the following command:

```
# hares -delete resource
```

  Note that deleting a resource won't take offline the object being monitored by the resource. The object remains online, outside the control and monitoring of VCS.

Adding, deleting, and modifying resource attributes

Resource names must be unique throughout the cluster and you cannot modify resource attributes defined by the system, such as the resource state.

To modify a new resource

- Type the following command:

```
# hares -modify resource attribute value
# hares -modify <resource> <attr> <value>
        [-sys <system>] [-wait [-time <waittime>]]
```

  The variable *value* depends on the type of attribute being created.

To set a new resource's enabled attribute to 1

- Type the following command:

```
# hares -modify resourceA Enabled 1
```

The agent managing the resource is started on a system when its Enabled attribute is set to 1 on that system. Specifically, the VCS engine begins to monitor the resource for faults. Agent monitoring is disabled if the Enabled attribute is reset to 0.

To add a resource attribute

```
# haattr -add resource_type  attribute [value]
      [dimension][default ...]
```

The variable *value* is a -string (default), -integer, or -boolean.

The variable *dimension* is -scalar (default), -keylist, -assoc, or -vector.

The variable *default* is the default value of the attribute and must be compatible with the *value* and *dimension*. Note that this may include more than one item, as indicated by ellipses (. . .).

To delete a resource attribute

```
# haattr -delete resource_type attribute
```

To add a static resource attribute

```
# haattr -add -static resource_type static_attribute [value]
    [dimension] [default ...]
```

To delete a static resource attribute

```
# haattr -delete -static resource_type static_attribute
```

To add a temporary resource attribute

```
# haattr -add -temp resource_type attribute [value]
    [dimension] [default ...]
```

To delete a temporary resource attribute

```
# haattr -delete -temp resource_type attribute
```

To modify the default value of a resource attribute

```
# haattr -default resource_type attribute new_value ...
```

The variable *new_value* refers to the attribute's new default value.

Defining attributes as local

Localizing an attribute means that the attribute has a per-system value for each system listed in the group's SystemList. These attributes are localized on a per-resource basis. For example, to localize the attribute *attribute_name* for *resource* only, type:

```
# hares -local resource  attribute_name
```

Note that global attributes cannot be modified with the `hares -local` command. The following table lists the commands to be used to localize attributes depending on their dimension.

Table 8-5 Making VCS attributes local

Dimension	Task and Command
scalar	Replace a value: **-modify [object]** ***attribute_name value*** **[-sys** ***system*]**
vector	■ Replace list of values: **-modify [object]** ***attribute_name value*** **[-sys** ***system*]** ■ Add list of values to existing list: **-modify [object] attribute_name -add value [-sys system]** ■ Update list with user-supplied values: **-modify [object]** ***attribute_name*** **-update** ***entry_value*** **... [-sys** ***system*]** ■ Delete all values in list (you cannot delete an individual element of a vector): **-modify [object]** ***attribute_name*** **-delete -keys [-sys** ***system*]**
keylist	■ Replace list of keys (duplicate keys not allowed): **-modify [object] attribute_name value ... [-sys system]** ■ Add keys to list (duplicate keys not allowed): **-modify [object]** ***attribute_name*** **-add** ***value*** **... [-sys** ***system*]** ■ Delete user-supplied keys from list: **-modify [object]** ***attribute_name*** **-delete** ***key*** **... [-sys** ***system*]** ■ Delete all keys from list: **-modify [object]** ***attribute_name*** **-delete -keys [-sys** ***system*]**

Table 8-5 Making VCS attributes local

Dimension	Task and Command
association	■ Replace list of key-value pairs (duplicate keys not allowed): **-modify [object]** ***attribute_name value*** **...** **[-sys** ***system*****]** ■ Add user-supplied list of key-value pairs to existing list (duplicate keys not allowed): **-modify [object]** ***attribute_name*** **-add** ***value*** **...[-sys** ***system*****]** ■ Replace value of each key with user-supplied value: **-modify [object]** ***attribute_name*** **-update** ***key value*** **...** **[-sys** ***system*****]** ■ Delete a key-value pair identified by user-supplied key: **-modify [object]** ***attribute_name*** **-delete** ***key*** **...** **[-sys** ***system*****]** ■ Delete all key-value pairs from association: **-modify [object]** ***attribute_name*** **-delete -keys** **[-sys** ***system*****]** **Note:** If multiple values are specified and if one is invalid, VCS returns an error for the invalid value, but continues to process the others. In the following example, if sysb is part of the attribute SystemList, but sysa is not, sysb is deleted and an error message is sent to the log regarding sysa. **hagrp -modify group1 SystemList -delete sysa sysb** **[-sys** ***system*****]**

Linking and unlinking resources

Link resources to specify a dependency between them. A resource can have an unlimited number of parents and children. When linking resources, the parent cannot be a resource whose Operations attribute is equal to None or OnOnly. Specifically, these are resources that cannot be brought online or taken offline by an agent (None), or can only be brought online by an agent (OnOnly).

Loop cycles are automatically prohibited by the VCS engine. You cannot specify a resource link between resources of different service groups.

To link resources

◆ Type the following command:

```
# hares -link parent_resource child_resource
```

The variable *parent_resource* depends on *child_resource* being online before going online itself. Conversely, *parent_resource* go offline before *child_resource* goes offline.

For example, a NIC resource must be available before an IP resource can go online, so for resources IP1 of type IP and NIC1 of type NIC, specify the dependency as:

```
# hares -link IP1 NIC1
```

To unlink resources

◆ Type the following command:

```
# hares -unlink parent_resource child_resource
```

Bringing resources online

To bring a resource online

◆ Type the following command:

```
# hares -online resource -sys system
```

Taking resources offline

To take a resource offline

◆ Type the following command:

```
# hares -offline [-ignoreparent|parentprop] resource -sys system
```

The option `-ignoreparent` enables a resource to be taken offline even if its parent resources in the service group are online. This option does not work if taking the resources offline violates the group dependency.

To take a resource and its parent resources offline

◆ Type the following command:

```
# hares -offline -parentprop resource -sys system
```

The command stops all parent resources in order before taking the specific resource offline.

!!! Editor: Need to verify command syntax.

To take a resource offline and propagate the command to its children

◆ Type the following command:

```
# hares -offprop [-ignoreparent] resource -sys system
```

As in the above command, the option `-ignoreparent` enables a resource to be taken offline even if its parent resources in the service group are online. This option does not work if taking the resources offline violates the group dependency.

Probing a resource

To prompt an agent to monitor a resource on a system

◆ Type the following command:

```
# hares -probe resource -sys system
```

Though the command may return immediately, the monitoring process may not be completed by the time the command returns.

Clearing a resource

To clear a resource

◆ Type the following command:

Initiate a state change from RESOURCE_FAULTED to RESOURCE_OFFLINE:

```
# hares -clear resource [-sys system]
```

Clearing a resource initiates the online process previously blocked while waiting for the resource to become clear. If *system* is not specified, the fault is cleared on each system in the service group's SystemList attribute.

See "To clear faulted, non-persistent resources in a service group" on page 263.

This command also clears the resource's parents. Persistent resources whose static attribute Operations is defined as None cannot be cleared with this command and must be physically attended to, such as replacing a raw disk. The agent then updates the status automatically.

Administering systems

To modify a system's attributes

◆ Type the following command:

```
# hasys -modify modify_options
```

Some attributes are internal to VCS and cannot be modified. For details on system attributes, see "About the -modify option" on page 236.

To display the value of a system's node ID as defined in the file /etc/llttab

◆ Type the following command:

```
# hasys -nodeid node_ID
```

To freeze a system (prevent groups from being brought online or switched on the system)

◆ Type the following command:

```
# hasys -freeze [-persistent] [-evacuate] system
```

-persistent	Enables the freeze to be "remembered" when the cluster is rebooted. Note that the cluster configuration must be in read/write mode and must be saved to disk (dumped) to enable the freeze to be remembered.
-evacuate	Fails over the system's active service groups to another system in the cluster before the freeze is enabled.

To unfreeze a frozen system (reenable online and switch of service groups)

◆ Type the following command:

```
# hasys -unfreeze [-persistent] system
```

Administering clusters

Retrieving version information

Retrieve information about the version of VCS running on the system.

```
# had -version
```

The command retrieves information about the engine version, the join version, the build date, and the PSTAMP.

```
# had -v
```

The command retrieves information about the engine version.

Administering resource types

Adding, deleting, and modifying resource types

After creating a resource type, use the command `haattr` to add its attributes. By default, resource type information is stored in the types.cf configuration file.

To add a resource type

```
# hatype -add resource_type
```

To delete a resource type

```
# hatype -delete resource_type
```

You must delete all resources of the type before deleting the resource type.

To add or modify resource types in main.cf without shutting down VCS

```
# hatype -modify resource_type SourceFile "./resource_type.cf"
```

The information regarding *resource_type* is stored in the file config/*resource_type*.cf, and an include line for *resource_type*.cf is added to the main.cf file. Make sure that the path to the SourceFile exists on all nodes before you run this command.

To set the value of static resource type attributes

```
# hatype -modify ...
```

Overriding resource type static attributes

You can override some resource type static attributes and assign them resource-specific values. When a static attribute is overriden and the configuration is saved, the main.cf file includes a line in the resource definition for the static attribute and its overriden value.

To override a type's static attribute

```
# hares -override resource static_attribute
```

To restore default settings to a type's static attribute

```
# hares -undo_override resource static_attribute
```

Using the -wait option in scripts

The `-wait` option is for use in scripts using VCS commands to change attribute values. The option blocks the VCS command until the value of the specified attribute is changed or until the timeout, if specified, expires. Specify the timeout in seconds.

The option can be used only with changes to scalar attributes.

The -wait option is supported with the following commands:

- **haclus**

  `haclus -wait` *`attribute value`* `[-clus` *`cluster`*`] [-time` *`timeout`*`]`

  Use the `-clus` option in a global cluster environment.

- **hagrp**

  `hagrp -wait group` *`attribute value`* `[-clus` *`cluster`*`] [-sys` *`system`*`] [-time` *`timeout`*`]`

  Use the `-sys` option when the scope of the attribute is local.

  Use the `-clus` option in a global cluster environment.

- **hares**

  `hares -wait resource` *`attribute value`* `[-clus` *`cluster`*`] [-sys` *`system`*`] [-time timeout]`

  Use the `-sys` option when the scope of the attribute is local.

  Use the -clus option in a global cluster environment.

- **hasys**

  `hasys -wait system` *`attribute value`* `[-clus` *`cluster`*`] [-time` *`timeout`*`]`

  Use the `-clus` option in a global cluster environment.

Administering simulated clusters from the command line

VCS Simulator is a tool to assist you in building and simulating cluster configurations. With VCS Simulator you can predict service group behavior during cluster or system faults, view state transitions, and designate and fine-tune various configuration parameters. This tool is especially useful when evaluating complex, multi-node configurations. It is convenient in that you can design a specific configuration without test clusters or changes to existing configurations.

You can also fine-tune values for attributes governing the rules of failover, such as Load and Capacity in a simulated environment. VCS Simulator enables you to simulate various configurations and provides the information you need to make the right choices. It also enables simulating global clusters.

See "Predicting VCS behavior using VCS Simulator" on page 405.

Configuring resources and applications in VCS

VCS detects the state of an application by continuously monitoring resources used by the application. If all resources required by the application are available, VCS declares the application available. VCS monitors resources using agents. See the *Veritas Cluster Server Bundled Agents Reference Guide* for a description of the agents provided by VCS.

Configuring resources and applications in VCS involves the following tasks:

- Creating a service group comprising all resources required for the application; configure resources.
  For example, to configure a database in VCS, you must configure resources for the database and for the underlying shared storage and network resources. Configuring a resource involves defining values for its attributes.
  The resources must be logically grouped in a service group. When a resource faults, the entire service group fails over to another node.
- Assigning dependencies between resources.
  For example, a MountV resource depends on a VMDg resource. Similarly, an IP resource depends on a NIC resource.
- Bringing the service group online.

VCS provides configuration wizards to configure commonly-used resources. You can also use Cluster Manager (Java Console), Cluster Management Console (Single Cluster Mode) also referred to as Web Console, or the command line to configure resources. When modifying agent attributes from the Java or Web Consoles, use a single forward slash (\) to denote path names. When editing the configuration file manually, use double forward slashes (\\).

Configuring shared storage

- If your configuration uses shared volumes or Logical Unit Numbers (LUNs) managed in a Network Appliance storage environment, use the NetAppSnapDrive and NetAppFiler agents.
- If your configuration uses shared disks and volumes managed using Veritas Storage Foundation for Windows (SFW), use the VMDg and MountV agents.

Before configuring shared storage, review the resource type and the attribute definitions of the VMDg and the MountV agents in the *Veritas Cluster Server Bundled Agents Reference Guide*.

For information about the NetAppFiler and NetAppSnapDrive agents, refer to the application-specific VCS implementation guide.

Managing storage using Network Appliance Filer

Network Appliance manages data by creating volumes on physical disks. These volumes can further be divided into Logical Unit Numbers (LUNs). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and Network Appliance SnapDrive installed. If you plan to use Fibre Channel (FC) for connecting the LUNs, ensure that you install the NetApp FCP Attach Kit on all the cluster nodes.

Note: Symantec does not support volumes created using qtree.

Perform the following tasks to create the required LUNs on the Network Appliance Filer and to make them accessible from cluster nodes:

- Create volumes on the Network Appliance Filer.
- Share the volumes.
- Create LUNs on the shared volumes.

Refer to Network Appliance documentation for instructions on performing these tasks.

Configuring Microsoft iSCSI Initiator

The Microsoft iSCSI initiator enables communication between Windows systems and Network Appliance Filers. The initiator uses the iSCSI protocol to present the filer volume as a local block device to the system.

To configure Microsoft iSCSI initiator

1. Make sure the Microsoft iSCSI Initiator software version 2.0 is installed on all cluster nodes. Refer to Microsoft documentation for further information.

2 Start the Microsoft iSCSI initiator. Double-click the Microsoft iSCSI Initiator icon from the desktop.

3 Click the **Target Portals** tab, if not already selected.

4 Click **Add...**.

5 In the Add Target Portals dialog box, specify the IP address or the DNS name for the Network Appliance Filer and click **OK**.

6 Click the **Available Targets** tab and click **Log On...**.

7 In the Log On to Target dialog box, verify the target portal name and select the **Automatically restore this connection when the system reboots** check box.

8 Click **OK**.

9 Click the **Persistent Target** tab to verify that the newly added target portal is listed under the **Select a target** box.

10 Click **OK**.

Connecting virtual disks to the cluster node

Once the required virtual disks are created on the Network Appliance Filer, they must be connected (if not connected already) to the cluster nodes using Network Appliance SnapDrive.

To connect virtual disks to the cluster node

1 Start the Computer Management MMC on the cluster node where you want to connect the LUN. Click **Start > All Programs > Administrative Tools > Computer Management**.

2 From the left pane, expand **Storage** and double-click **SnapDrive**.

3 Double-click **Disks** to see the LUNs that are connected to the node.

4 Right-click the LUN you want to connect and then click **Connect Disk...**.

5 In the Connect Disk alert box, click **OK**.

Disconnecting virtual disks from the cluster nodes

Steps to disconnect the virtual disks from a cluster node.

To disconnect virtual disks

1 Start the Computer Management MMC on the cluster node where you want to disconnect the LUN. Click **Start > All Programs > Administrative Tools > Computer Management**.

2 From the left pane, expand **Storage** and double-click **SnapDrive**.

3 Double-click **Disks** to see the LUNs that are connected to the node.

4 Right-click the LUN you want to disconnect and then click **Disconnect Disk...**.

5 In the Disconnect Disk alert box, click **OK**.

Using SFW with VCS

The following advanced features of SFW require special consideration when used in a VCS environment. Review the Storage Foundation documentation for more information.

- Deporting Disk Groups
  SFW does not allow disk groups configured as VCS resources to be deported. They must be brought online or taken offline using VCS.
- Dynamic Group Split and Join (DGSJ)
  SFW does not allow splitting a disk group configured as a VCS resource if the split operation causes a volume configured as a VCS resource to be part of the target group.
  SFW does not allow a disk group configured as a VCS resource to be the source disk group in a join operation.
- Deleting Volumes
  SFW does not allow deleting volumes configured as VCS resources.
- Volume Snap Back
  If a volume formed as a result of a Prepare and Snap Shot operation, is configured as a VCS resource, SFW does not allow Snap Back operations on the volume. (See the Veritas Storage Foundation documentation for more information about these operations.)

Prerequisites

- Verify that SFW HA or VCS for NetApp SnapMirror is installed on all cluster systems.
- Configure the clustered disk group using Storage Foundation. Verify the disk group contains shared disks only.
- Disable the option **Reset SCSI Bus at IC Initialization** from the SCSI Select utility.
- Create a separate clustered disk group for each application to be clustered. Do not create a clustered disk group for more than one application. Configure all volumes or LUNs to be part of the VCS configuration and of the same service group.
- Assign a unique disk group name to each clustered disk group within a cluster.
- Ensure that the device path to the shared disk group or LUNs is recognized by all systems sharing the disk.

Configuration tasks

- In your service group, create the following resources:
  - For SFW HA, create resources of type VMDg and MountV.
  - For VCS for NetApp SnapMirror, create resources of type NetAppFiler and NetAppSnapDrive.

  See "Adding a resource" on page 200 for instructions.
- Configure the following required attributes for these resources:

  VMDg Resource
  - DiskGroupName: The name of the cluster disk group. Retrieve the name by running the command `vxdg list`, or by using the VMGetDrive utility.

    See "The vmgetdrive utility" on page 680 for more information.

  MountV Resource
  - MountPath: The drive letter or path to an empty NTFS folder that will be assigned to the volume being mounted.
  - VolumeName: The name of the volume to be mounted. For example, the name could be Raid1, Stripe2, Volume01, etc. Use the VMGetDrive utility to retrieve the volume name.

    See "The havol utility" on page 676 for instructions.
  - VMDGResName: The name of the Volume Manager Diskgroup (VMDg) resource on which the MountV resource depends.

NetAppFiler resource

- FilerName: DNS-resolvable name or IP address of the locally attached filer.
- StorageIP: The private storage IP address of the filer.

NetAppSnapDrive resource

- FilerResName: Name of the VCS NetAppFiler-type resource in the service group.
- VolumeName: Name of the volume containing the virtual disk. Define the volume name in the same case as on the filer.
- ShareName: Name of the CIFS share containing the virtual disk.
- LUN: Name of the LUN on the filer that is presented to the host for mounting. Define the LUN name in the same case as on the filer.
- MountPath: Drive letter to be assigned to the virtual disk.
- Initiator: Name of the iSCSI or FC initiator that the host uses to connect to the virtual disks on the filer. You can retrieve this from the Disk Management Console.

- Link the resources as follows:
  - For SFW HA, link MountV and VMDg resources such that the MountV resource depends on the VMDg resource.
  - For VCS for NetApp SnapMirror, link NetAppSnapDrive and NetAppFiler resources such that the NetAppSnapDrive resource depends on the NetAppFiler resource.

  See“Linking resources” on page 213 for instructions.
- Configure other resources in the service group, if required.
- Bring the MountV or the NetAppSnapDrive resource online.

Configuring network resources

This section provides an overview of the steps involved in configuring network resources in a VCS cluster.

- If each system in your cluster uses a single adapter for the public network, use the NIC and IP agents.
- If your cluster systems use multiple adapters for the public network, use the IPMultiNICPlus agent.
- If your cluster systems use virtual computer names, use the Lanman agent.

Configuring IP addresses in single-adapter systems

Before configuring the network resources, review the resource type and the attribute definitions of the NIC and IP agents, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Prerequisites

- Ensure that the NIC has the correct administrative IP address and subnet mask.
- If the NICs have built-in failover support, disable it. Refer to the documentation provided by the NIC vendor.
- Do not configure IP addresses added from the Control Panel.
- Verify that the virtual IP address to be assigned is unique and does not exist in the network.
- Disable DHCP on the NIC.

To disable DHCP

1. Open the Network Connections Control Panel.
2. Right-click the network connection and click **Properties**.
3. In the Properties dialog box for the respective local area connection, select the **General** tab, if not already selected.
4. Select **Internet Protocol (TCP/IP)** and click **Properties**.
5. Verify that the **Obtain an IP address automatically** option is *not* selected.
6. Specify values for **IP address**, **Subnet mask**, and **Default Gateway**, if not already specified.
7. Click **OK** and close the Control Panel.

Configuration tasks

1 In your service group, create resources of type NIC and IP.
See "Adding a resource" on page 200 for instructions.

2 Configure the following required attributes for these resources:
NIC Resource
   - MACAddress: The physical address of the NIC to be monitored. You can retrieve the physical addresses of NICs using the command `ipconfig -all`. Note that this attribute is always local.
   - UseConnectionStatus: A flag that defines whether the NIC maintains its connection status.

   IP Resource
   - Address: The unique virtual IP address to be assigned to the NIC.
   - MACAddress: The physical address of the NIC to which the virtual IP address is assigned. Note that this attribute is always local.
   - SubNetMask: The subnet mask associated with the IP address.

   Ensure that the value of the attribute UseConnectionStatus is correct. This value is set to True by default, and indicates that all NICs maintain their connection status. If UseConnectionStatus is set to False, ensure that the NIC has an IP address assigned and that at least one host is listed in the attribute PingHostList.

3 Link the IP and NIC resources such that the IP resource depends on the NIC resource.
See "Linking resources" on page 213 for instructions.

4 Configure other resources in the service group, if required.

5 Bring the IP resource online.

Configuring IP addresses in multiple-adapter systems

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

Prerequisites

- Ensure that the NIC has the correct administrative IP address and subnet mask.
- Disable DHCP.
  See "Configuring IP addresses in single-adapter systems" on page 283 for instructions.
- If the NICs have built-in failover support, disable it. Refer to the documentation provided by the NIC vendor.
- Do not configure IP addresses added from the Control Panel.
- Verify that the virtual IP address to be assigned is unique and does not exist in the network.
- Verify that NICs on a single system are connected to the same network segment.

Configuration tasks

1. In your service group, create a resource of type IPMultiNICPlus.
   See "Adding a resource" on page 200 for instructions.
2. Configure the following required attributes for the resource:
   - AdapterList: A list of MAC addresses (physical addresses) of NICs that form the IPMultiNicPlus resource.
   - Address: The unique virtual IP address to be assigned to the active NIC.
   - AdminIPAddr: Unique administrative IP address assigned to the active NIC in the AdapterList attribute. Note that this attribute is always local. If this attribute is configured incorrectly, the system will not be accessible in the network until you bring the IPMultiNicPlus resource online on the system.
   - AdminSubnetMask: The subnet mask associated with the administrative IP address. Note that this attribute is always local.
   - SubNetMask: The subnet mask associated with the virtual IP address.
3. Configure other resources in the service group, if required.

4 Bring the IPMultiNICPlus resource, and other resources in the service group, online.

Configuring virtual computer names

Before configuring the agent, review the resource type and the attribute definitions of the Lanman agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Prerequisites

- Remove static entries mapping the virtual name to the IP address from your WINS server.
- If using the agent to bind multiple IP addresses to a virtual computer name, make sure the IP addresses belong to different subnets.
- Make sure the VCS Helper domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This will enable the Lanman agent to refresh the resource records on the DNS servers. See the Lanman agent description in the *Veritas Cluster Bundled Agents Reference Guide*.

Configuration tasks

1 In your service group, create resources of type NIC and IP.
See "Configuring IP addresses in single-adapter systems" on page 283 for instructions.

2 Create a resource of type Lanman.

3 Configure the following required attributes for the resource:
   - VirtualName: The virtual computer name to be assigned to the server.
   - IPResName: The name of the IP resource on which the Lanman resource depends. The IPResName attribute is not required if you have the MultiNet attribute set to 1.

4 Link the IP and NIC resources such that
   - the IP resource depends on the NIC resource, and
   - the Lanman resource depends on the IP resource.

See "Linking resources" on page 213 for instructions.

5 Configure other resources in the service group, if required.

6 Bring the Lanman resource, and other resources in the service group, online.

Configuring file shares

This section describes how to configure file shares in VCS.

- To configure a shared directory, use the FileShare agent.
- To configure multiple directories, use the Composite FileShare agent.

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

Creating a file share service group using the wizard

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

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

Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Verify that the directories to be shared reside on shared drives.
- Mount the drives or LUNs containing the shared directories from the system from which you run the wizard. Unmount the drives or LUNs from other systems in the cluster.
- Verify that Veritas Command Server is running on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
  - A unique virtual computer name to be assigned to the file share server. This is the name by which clients will access the server. The virtual name must not exceed 15 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
  - A unique virtual IP address to be assigned to the file share server.
  - The list of directories to be shared.

The wizard enables you to add existing shares to the VCS configuration. However, you cannot add special shares (shares created by the operating system for administrative and system use). For example, you cannot add the shares ADMIN$, print$, IPC$, and *DriveLetter*$ to the VCS configuration.

Configuration tasks

To configure a file share

1 Start the File Share Configuration Wizard. **(Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > File Share Configuration Wizard**)

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

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

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

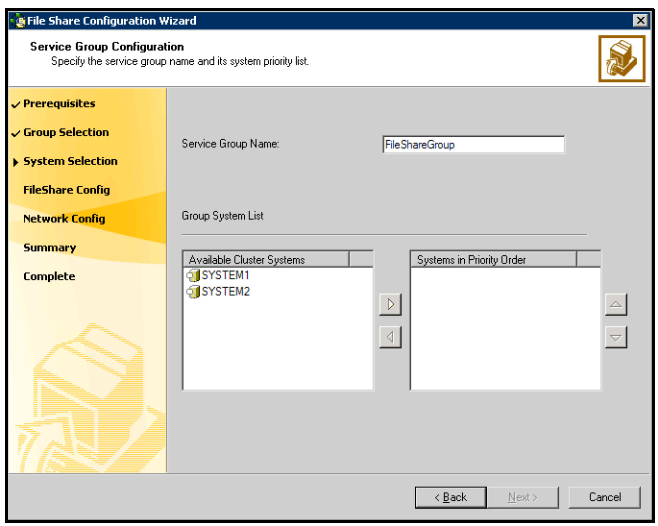

Service Group Name	Type a name for the file share service group.

Available Cluster Systems	Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list. To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow. To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

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

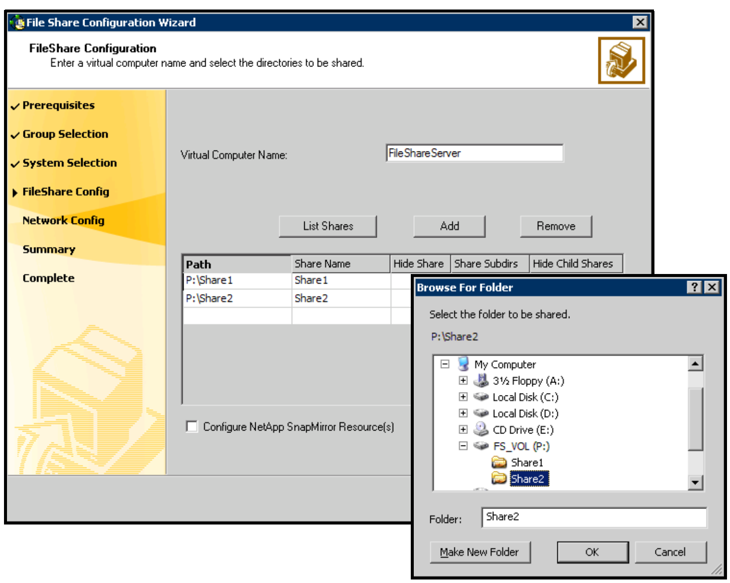

Virtual Computer Name	Type a unique virtual computer name by which the server will be known to clients. The virtual name must not exceed 15 characters.

List Shares	Click **List Shares** to view the existing shares on the shared storage, then select a share and click **Add**.
Add	Click **Add** to add a file share.
Path	Type the path of the directories to be shared or click the field and then click the ellipse icon (...) to browse for folders. The selected directories must meet the following conditions: ■ The selected drive, the mount path, and the file path must not exist in the VCS configuration. ■ The directories to be shared must reside on shared, non-system drives. The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.
Share Name	If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can select a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.
Hide Share	Check **Hide Share** check box to make the new share a hidden share.
Share Subdirs	Check the **Share Subdirs** check box to share the subdirectories.
Hide Child Shares	Check the **Hide Child Shares** check box to hide the shared subdirectories.
Remove	To remove a file share from the configuration, click to select the file share, and then click **Remove**.
Configure NetApp SnapMirror Resource(s)	This is applicable in case of VCS for NetApp SnapMirror only. Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

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

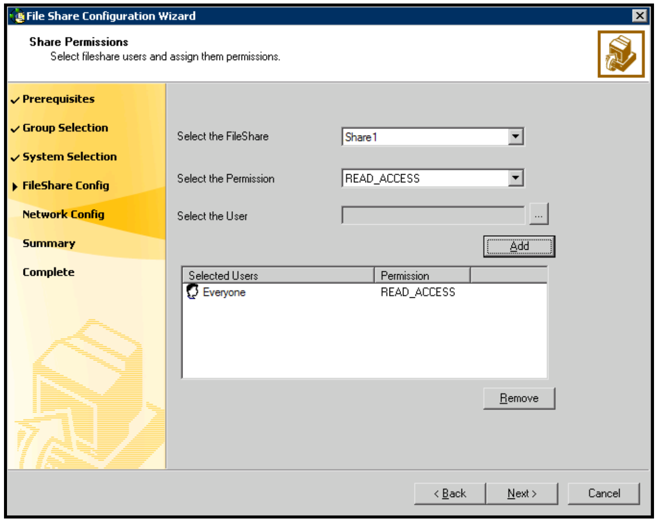

Select the FileShare	From the drop-down list, select the file share with which to associate user permissions, or select the default **All FileShares** to set the same permissions for all file shares.
Select the Permission	From the drop-down list, select the permission to be associated with the user.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

7 This is applicable in case of VCS for NetApp SnapMirror only.

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

If you are configuring MPIO over FC, you must select at least 2 FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

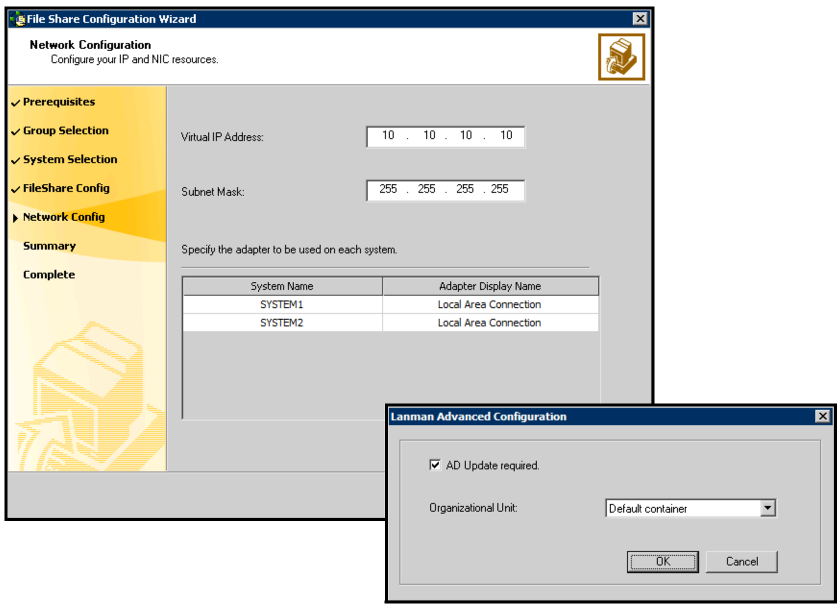

Virtual IP Address	Type a unique virtual IP address for the virtual server.
Subnet Mask	Type the subnet to which the virtual server belongs.

Advanced Settings	Click **Advanced Settings...** to specify additional details for the Lanman resource. On the Lanman Advanced Configuration dialog box, complete the following: **1** Check **AD Update required** check box to enable the Lanman resource to update the Active Directory with the virtual name. **2** From the Organizational Unit drop-down list, select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." **3** Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
Adapter Display Name	Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.

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

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

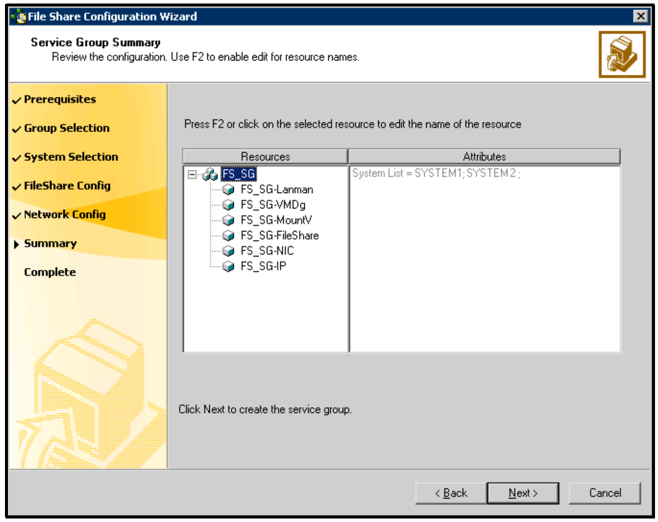

Resources	Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required. To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes.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

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

Modifying a file share service group using the wizard

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

- If the file share service group is online, you must run the wizard from a system on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.

- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), and NetAppSnapDrive and NetAppFiler (in case of VCS for NetApp SnapMirror) resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

To modify a file share service group

1. Start the File Share Configuration Wizard on a system on which the file share service group is online. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > File Share Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. On the Wizard Options panel, click **Modify service group**, select the service group to be modified, and click **Next**.
4. Follow the wizard instructions described between step 4 through step 10 in "Creating a file share service group using the wizard" on page 288.

Deleting a file share service group using the wizard

This section describes steps to delete a file share service group using the configuration wizard.

To delete a file share service group

1. Start the File Share Configuration Wizard on a system configured to host the file share service group. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > File Share Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. On the Wizard Options panel, click **Delete service group**, select the service group to be deleted, and click **Next**.
4. On the Service Group Summary panel, click **Next**. A message appears informing you that the wizard will run commands to delete the service group. Click **Yes** to delete the service group.
5. Click **Finish**.

Configuring multiple file shares

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

Prerequisites

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

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

Configuration tasks

1. Configure your shared storage.
   See "Configuring shared storage" on page 278.
2. Configure the NIC and IP resources.
   See "Configuring IP addresses in single-adapter systems" on page 283.
3. Create a resource of type CompositeFileShare.
4. Configure the following required attributes for the resource:
   - MountResName: The name of the MountV resource on which the Composite FileShare resource depends.
   - PathAndShareName: A list specifying the respective paths and share names of the directories to be shared. If the path of a shared directory is \Documents, and the share name is UserDocs, the attribute is defined in the configuration file as PathandShareName is {"\\Documents" = "UserDocs"}.

   To create a hidden share, set the HiddenShare attribute to 1. Do not append the share name with a $ (dollar) sign.
5. Configure a Lanman resource. Do not create other resources on which the Lanman agent depends.
   See "Configuring virtual computer names" on page 286 for instructions.

6 Link resources to create the following dependencies:
   - CompositeFileShare resource depends on the MountV (in case of SFW HA) or NetAppSnapDrive (in case of VCS for NetApp SnapMirror) resource.
   - Composite FileShare resources depends on the Lanman resource.
   - Lanman resource depends on IP resource, which in turn depends on the NIC resource.

   See "Linking resources" on page 213 for instructions.
7 Configure other resources in the service group, if required.
8 Bring the Lanman resource, and other resources in the service group, online.

Configuring print shares

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

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

Creating a print share service group using the wizard

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

See "Modifying a print share service group using the wizard" on page 311 for instructions on modifying a print share service group.

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

Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Verify that VCS Command Server is running on all systems in the cluster.
- Verify that the network printer has an IP address assigned.
- Symantec recommends creating spooler and the replication directories on different disk partitions or volumes or LUNs.
- Mount the drives or LUNs with the spooler and the replication directories from the system on which you run the wizard. Unmount the drives or LUNs from other systems in the cluster.
- Install software drivers for the network printer on all systems in the cluster. See "Installing printer drivers" on page 301 for instructions.
- Verify that you have the following information ready. The wizard will prompt you for this information:
  - A unique virtual computer name to be assigned to the print share server.

This is the name by which clients will access the server. The virtual name must not exceed 15 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.

- A unique virtual IP address to be assigned to the print share server.
- The network printer's IP address.

Installing printer drivers

Add printers as per the following instructions.

To add a print driver (for Windows Server 2003)

Run the Print Share Configuration wizard to add printers to the cluster. The drivers will be installed on the shared disk in the service group and moved from server to server along with the service group.

Configuration tasks

Creating a print share service group involves the following tasks:

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

To create a print share service group with a PrintSpool resource

1. Start the Print Share Configuration Wizard. **(Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Print Share Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. On the Wizard Options panel, click **Create service group** and click **Next**.

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

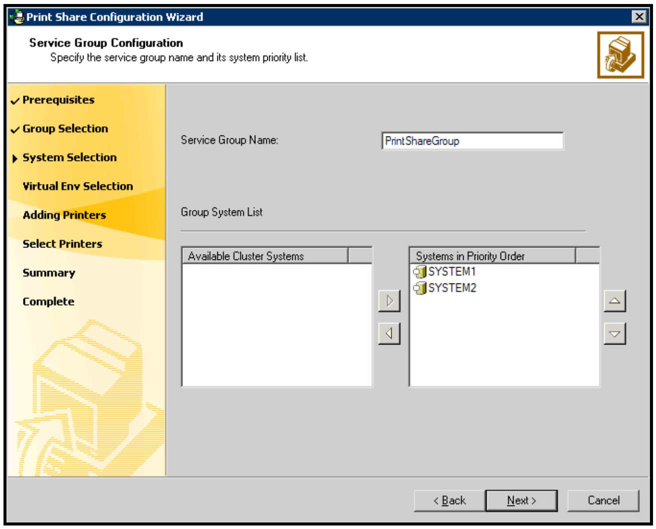

Service Group Name	Type a name for the print share service group.
Available Cluster Systems	Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list. To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow. To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

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

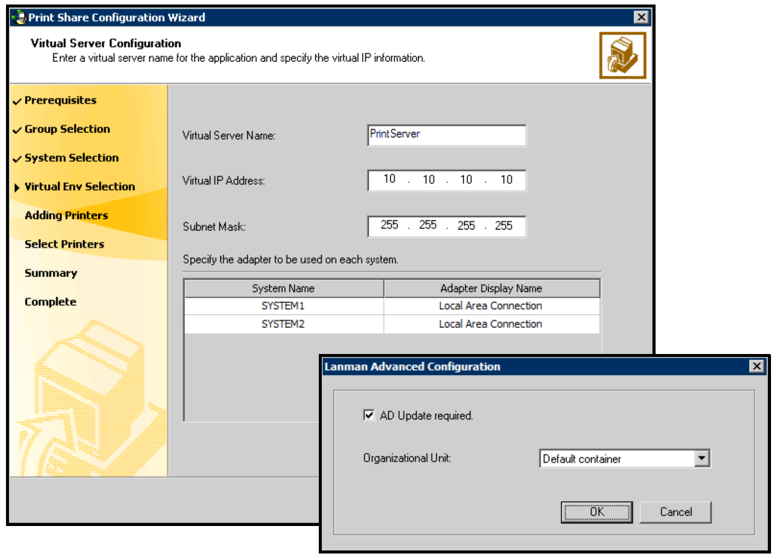

Virtual Server Name	Type a unique virtual computer name by which the server will be known to clients. Note that the virtual name must not exceed 15 characters.
Virtual IP Address	Type a unique virtual IP address for the virtual server.
Subnet Mask	Type the subnet to which the virtual server belongs.

Advanced Settings	Click **Advanced Settings...** to specify additional details for the Lanman resource. On the Lanman Advanced Configuration dialog box, complete the following: 1 Check **AD Update required** check box to enable the Lanman resource to update the Active Directory with the virtual name. 2 From the Organizational Unit drop-down list, select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers." 3 Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
Adapter Display Name	Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.

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

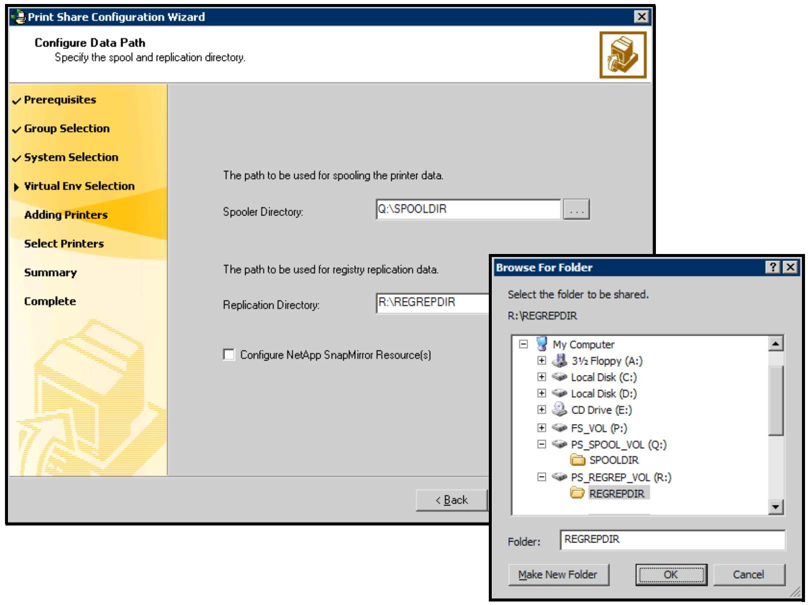

Spooler Directory

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

Replication Directory

Type the path or click **...** (ellipsis button) to browse for the directory. All changes related to the printer registry keys will be logged at this location.
The selected directories must fulfill the following conditions:

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

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

7 This is applicable in case of VCS for NetApp SnapMirror only.

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

If you are configuring MPIO over FC, you must select at least 2 FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

PrintSpool resource depends, including the Lanman and ServiceMonitor resources.

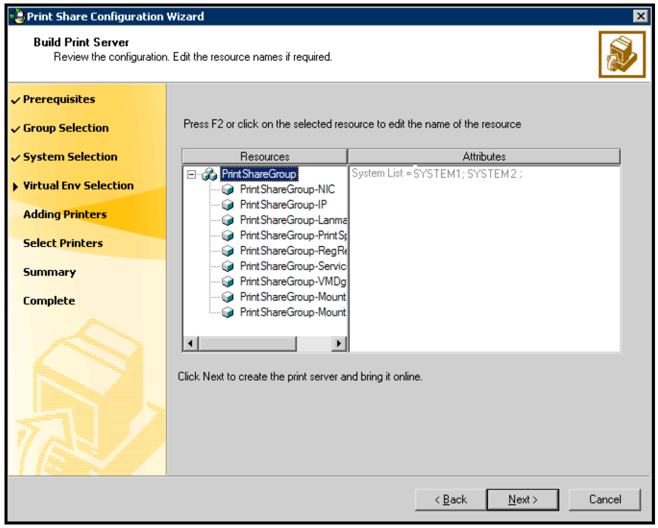

Resources	Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required. To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes.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

9 Bring the PrintSpool resource online.

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

To add the network printer to the virtual computer

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

To launch the Add Printer wizard, return to the Print Share Configuration Wizard and click **Add Printer** on the Add Printer panel, or in Windows Explorer, search for the virtual computer, explore the virtual computer by double-clicking its name and on the virtual computer's Printers folder, double-click **Add Printer**.

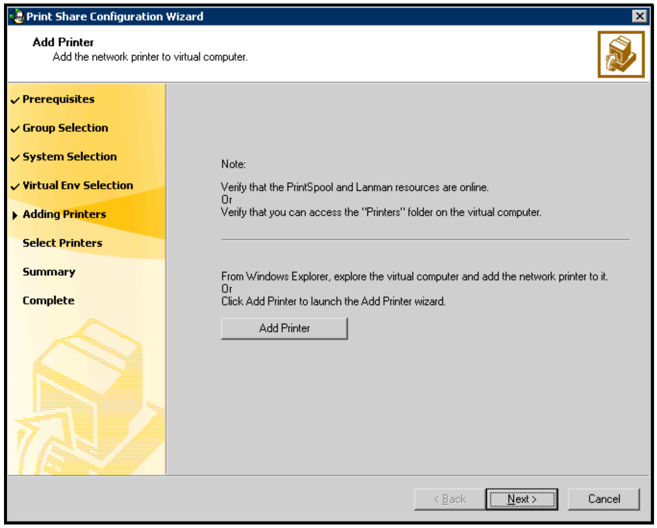

2 In the Add Printer wizard, review the information in the Welcome panel and click **Next**.

3 Follow the wizard instructions to add the network printer to the virtual computer.
In the Printer Sharing dialog box, always choose the **Do not share this printer** option.
Repeat these steps for each additional printer to be installed.

4 Return to the Print Share Configuration Wizard, and proceed to the next step to configure a PrintShare resource in your service group and bring it online.

To configure a PrintShare resource for the service group

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

2 On the Printer List panel, specify the printers to be included in the print share service group and then click **Next**.

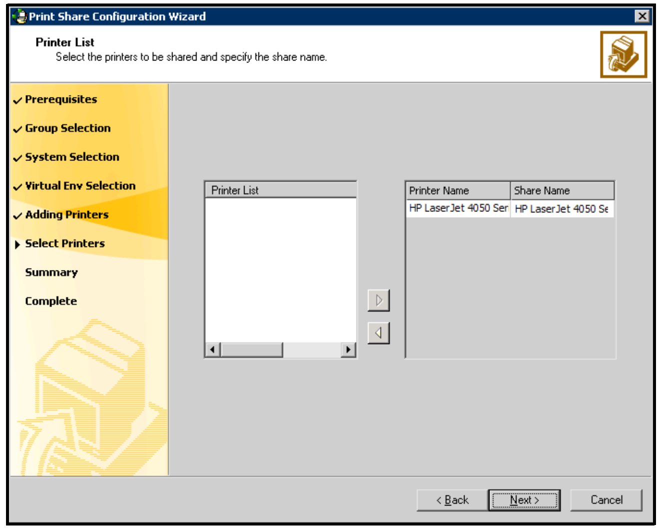

Printer List	Click to select the printer, and then click the right arrow to include the selected printers in your service group. To remove a selected printer from your service group, click the printer from the Printer Name list and click the left arrow.
Share Name	Type a unique share name for the printer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

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

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

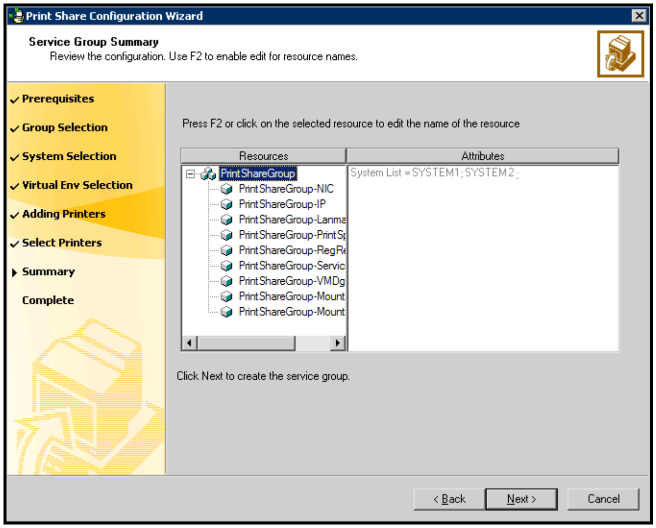

Resources	Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required. To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes.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

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

Modifying a print share service group using the wizard

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

- If the print share service group is online, you must run the wizard from a system on which the service group is online. You can then add and remove resources to the configuration using the wizard; you cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), and NetAppSnapDrive and NetAppFiler (in case of VCS for NetApp SnapMirror) resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

To modify the print share service group

1. Start the Print Share Configuration Wizard on a system on which the print share service group is online. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Print Share Configuration Wizard**)
2. Read the information on the Welcome panel and click **Next**.
3. On the Wizard Options panel, click **Modify service group**, select the service group to be modified, and click **Next**.
4. Follow the wizard instructions and make desired modifications to the service group configuration. See "Creating a print share service group using the wizard" on page 300 for more information about the configuration wizard.
   If you are modifying the service group to remove a PrintShare resource, make sure you offline the resource before deleting it.

Deleting a print share service group using the wizard

This section describes steps to delete a print share service group using the configuration wizard.

To delete the print share service group

1. Start the Print Share Configuration Wizard on a system configured to host the print share service group. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Print Share Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Delete service group**, select the service group to be deleted, and click **Next**.
4. In the Service Group Summary panel, click **Next**.
5. A message appears informing you that the wizard will run commands to delete the service group. Click **Yes** to delete the service group.
6. Click **Finish**.

Configuring Microsoft virtual machines

The MSVirtualMachine agent provides high availability to virtual machines created using Microsoft Virtual Server on Windows 2003 systems. Before configuring a virtual machine, review the resource types and attribute definitions of the MSVirtualMachine agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Before configuring Microsoft virtual machines in VCS

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

Configuring the virtual machine service group using the wizard

The Microsoft Virtual Machine Configuration Wizard creates and modifies service groups for Microsoft virtual machines, making virtual machines highly available.

The wizard duplicates the virtual machine and its network configuration on other nodes in the cluster. If the wizard discovers a virtual machine with the same name on another node, the wizard does not duplicate the virtual machine. The wizard does not duplicate virtual hard disks.

You must create a different service group for each virtual machine.

Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Disable the firewall on each node that will host the service group.
- You must have the following information ready. The wizard will prompt you for this information:
  - The name of the virtual machine.
  - Destination on shared disks or LUNs for the virtual hard disk files.
  - Network adapters on physical nodes to be associated with network adapters on the virtual machine.
  - Information about monitoring heartbeats (optional).

Configuration tasks

1. Start the Microsoft Virtual Machine Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Microsoft Virtual Machine Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Create service group** and click **Next**.
4. Enter a name for the service group and specify the systems on which to configure the service group.
   - Enter a name for the service group.
   - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.
     To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.
   - To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
   - Click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.

5 On the Microsoft Virtual Machine Group Configuration panel, specify details about the virtual machine and then click **Next**.

- Select the virtual machine.
- For each virtual disk, specify a destination folder where the virtual hard disk files will be moved. Click the Browse icon to browse for folders.
- To enable detail monitoring for the virtual machine, check the **Monitor Heartbeats** check box and
- In the **No. of Monitor Cycles** field, type failed heartbeat threshold. The threshold defines the number of consecutive monitor cycles the agent waits to detect heartbeats from the virtual machine before declaring the resource as faulted.
- Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. This is applicable in case of VCS for NetApp SnapMirror only.

  The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration.

  Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

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

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

  The fields for the virtual IP address and subnet mask are disabled by design.
- Click **Next**.

7 This is applicable in case of VCS for NetApp SnapMirror only.

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

If you are configuring MPIO over FC, you must select at least 2 FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

8 Review the service group configuration.

- The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.
- The wizard assigns unique names to resources. Change names of resource, if required.

  To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.
- Click **Next**.
- A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.

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

9 In the completion dialog box, select the check box if you want to bring the service group online on the local system.

10 Click **Finish**.

Note: For this release, you must configure the virtual network associated with the virtual machine manually on each node in the cluster. Do not switch the service group before configuring the network.

Configuring IIS sites

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

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

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

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

Before configuring IIS sites

- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify the sites to be monitored are on shared storage.
- Do not use the IIS agent to configure SMTP and NNTP sites if you have Microsoft Exchange installed.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- Verify the port numbers assigned to IIS sites are not used by other programs.
- Synchronize the IIS configuration on all nodes hosting the service group, as instructed in the next section.

To synchronize the IIS configuration on Windows 2003 systems

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

1. Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system.

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

   ```
   %systemroot%\System32> iiscnfg /copy /ts target_system /tu
   user_name /tp password
   ```
2. Stop and restart IIS Admin Service on all nodes.

Creating an IIS service group using the wizard

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

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

Prerequisites

- Verify that you have Administrator privileges on the system from where you run the wizard.
- Verify that the VCS engine, HAD, is running on the system from which you run the wizard.
- Mount the drives (in case of SFW HA), or connect the LUNs (in case of VCS for NetApp SnapMirror) containing the shared directories on the system from which you run the wizard. Unmount the drives or LUNs from other systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
  - IIS sites to be monitored.
  - Application pools associated with each site.
  - Port numbers associated with each site.
  - Virtual IP addresses and computer names associated with the sites. The virtual IP addresses and the virtual computer names must have forward and reverse entries in the DNS.

Configuration tasks

1. Start the IIS Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > IIS Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. On the Wizard Options panel, click **Create service group** and click **Next**.

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

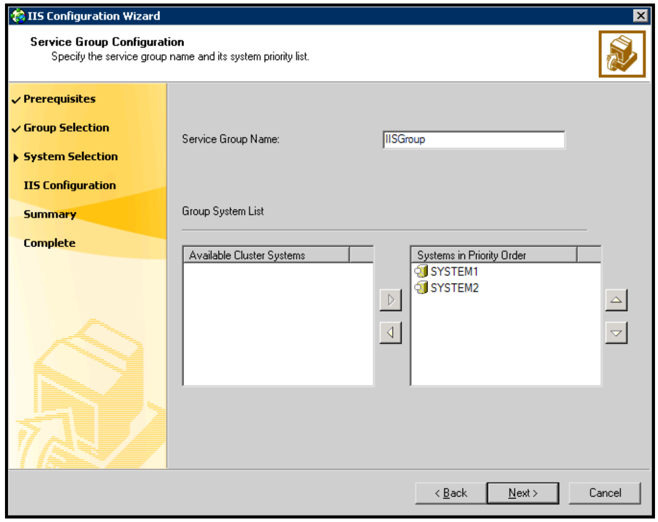

Service Group Name	Type a name for the IIS service group.
Available Cluster Systems	Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list. To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow. To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

5 On the Configure IIS Sites panel, add and remove sites from the service group, configure IP addresses, ports, and virtual computer names,

optionally choose to configure NetApp SnapMirror resources, and then click **Next**.

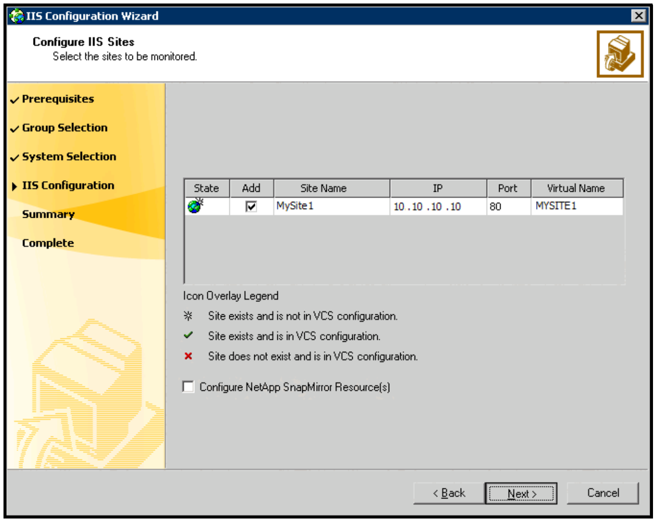

Add	Check the check box corresponding to the site to be configured in VCS.
IP	Type the virtual IP address for each site to be configured. Make sure that each virtual IP address is associated with only one virtual computer name and vice-versa.
Port	Type the port number for each site to be configured.
Configure NetApp SnapMirror Resource(s)	This is applicable in case of VCS for NetApp SnapMirror only. Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

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

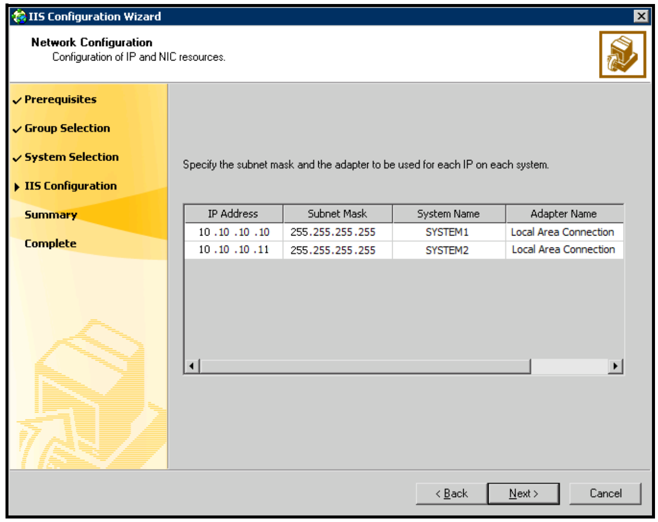

IP Address	Displays the virtual IP addresses. The wizard groups systems by the virtual IP addresses associated with the systems.
Subnet Mask	Type the subnet mask associated with each virtual IP address.
Adapter Name	Select the adapter associated with the virtual IP address on each system.

7 This is applicable in case of VCS for NetApp SnapMirror only.
On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.
If you are configuring MPIO over FC, you must select at least 2 FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

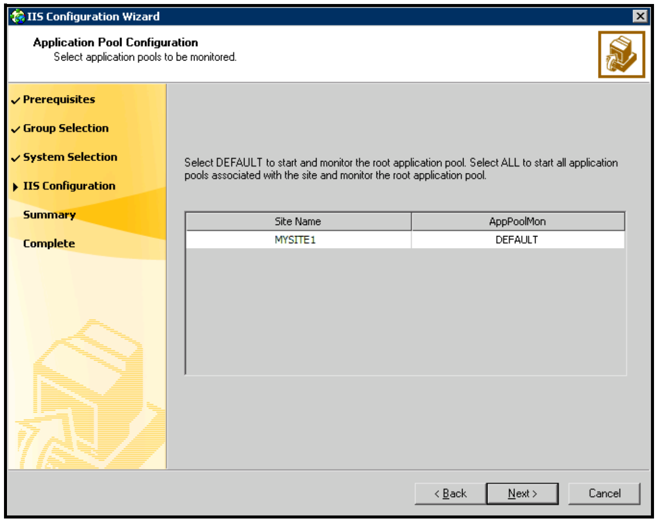

Site Name	Displays the site names.
AppPoolMon	For each site, select the monitoring options from the AppPoolMon list. **NONE**—The agent will not monitor the application pool associated with the site. **DEFAULT**—Starts and monitors the root application pool associated with the site. **ALL**—Starts all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

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

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

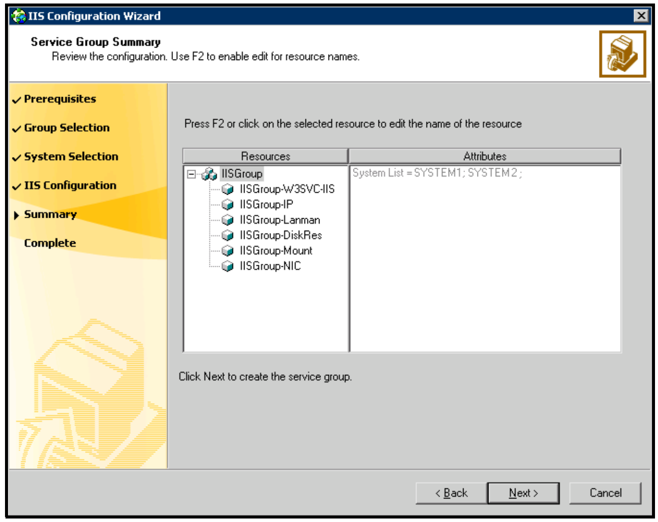

Resources	Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required. To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes.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

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

Modifying an IIS service group using the wizard

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

- If the IIS service group is online, you must run the wizard from a system on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), and NetAppSnapDrive and NetAppFiler (in case of VCS for NetApp SnapMirror) resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

To modify the IIS service group

1. Start the IIS Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > IIS Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Modify service group**, select the service group to be modified, and click **Next**.
4. Follow the wizard instructions and make desired modifications to the service group configuration. See "Creating an IIS service group using the wizard" on page 318 for more information about the configuration wizard.

Deleting an IIS service group using the wizard

This section describes steps to delete an IIS service group using the configuration wizard.

To delete the IIS service group

1. Start the IIS Configuration Wizard on a system configured to host the IIS service group. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > IIS Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Delete service group**, select the service group to be deleted, and click **Next**.
4. In the Service Group Summary panel, click **Next**. A message appears informing you that the wizard will run commands to delete the service group. Click **Yes** to delete the service group.
5. Click **Finish**.

Configuring services

This section provides an overview of the steps involved in configuring services in a VCS cluster.

- To start, stop, and monitor a service, use the GenericService agent.
- To monitor a service, use the ServiceMonitor agent.

Configuring a service using the GenericService agent

The GenericService agent starts, stops, and monitors services. Before configuring the service group, review the resource types and attribute definitions of the GenericService agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

You can configure the GenericService agent manually, as described below, or by using the Application Configuration Wizard.

See "Configuring applications using the Application Configuration Wizard" on page 343 for instructions.

Prerequisites

- Change the startup type of the generic service to **Manual**.
- If monitoring the service in a user-context, configure the service to start in the context of the specified user account. Make sure the check box for **Allow service to interact with desktop** is unchecked.

To change a service startup type to Manual

1. Open the Services Control Manager.
2. Right-click the service and click **Properties**.
3. In the **Properties** dialog box, click the **General** tab.
4. From the **Startup Type** list, select **Manual**.
5. Click **OK**.
6. Close the **Services** Control Manager.

To configure a service to start in a user-context

1. Open the Services Control Manager.
2. Right-click the service and click **Properties**.
3. In the **Properties** dialog box, click the **LogOn** tab.

4 Click **This Account**.

5 Click **Browse** to browse existing user accounts.

6 In the **Select User** dialog box, click the user in whose context you want to run the service and click **OK**.

7 Enter the password for the selected user.

8 Click **OK** and close the **Services** Control Manager.

Configuration tasks

1 In your service group, create a resource of type GenericService.
See "Adding a resource" on page 200 for instructions.

2 Configure the following required attribute for the resource.

- ServiceName: The name of the service to be monitored, as displayed by SCM.

If required, configure the following optional attributes for the resource:

- UserAccount: A valid user account in whose context the service will be monitored. Username can be of the form *username@domain.com* or *domain.com\username*. If you do not specify a value for this attribute, then the user account of the service in the SCM is ignored. To monitor service under built-in accounts, you must provide explicit values. For example:
  - On Windows 2003: User Account="LocalSystem", "Local Service", or "Network Service". Domain="NT Authority".
- Password: The password for the user account.
- Domain: The domain name to which the user specified in the UserAccount attribute belongs.

3 Configure other resources in the service group, if required.

4 Bring the GenericService resource, and other resources in the service group, online.

Configuring a service using the ServiceMonitor agent

The ServiceMonitor agent monitors a service or starts a script that monitors a service. Before configuring the service group, review the resource types and attribute definitions of the agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

You can configure the agent manually, as described below, or by using the Application Configuration Wizard, described in "Configuring applications using the Application Configuration Wizard" on page 343.

Prerequisites

- If using the agent to start a script, copy the script a locally on each system in the cluster.
- If using the agent to monitor a service, start the service in the context of the LocalSystem account or in the context of the user account specified in the configuration.
- Verify that the user in whose context the service or script needs to be started, exists as a domain user or LocalSystem user.

Configuration tasks

1. In your service group, create a resource of type ServiceMonitor.
   See "Adding a resource" on page 200 for instructions.
2. Configure the following required attribute for the resource.
   - ServiceorScriptName: The name of the service to be monitored using the Service Control Manager (SCM). When monitoring the service through a user defined script, specify the complete path of the script, including any command-line arguments.

   When monitoring a service through a user-defined script, values for the following attributes also need to be specified:
   - MonitorService: A flag that defines whether the agent monitors a service using the SCM or starts a script to monitor a service. If the flag is set to 1, the agent monitors a service specified by the attribute ServiceOrScriptName. If the flag is set to 0 the agent starts a script specified by the attribute ServiceOrScriptName. Default is 1.
   - MonitorProgTimeout: The maximum wait time, in seconds, for the agent to receive a return value from the monitor script. This attribute is ignored if the MonitorService flag is set to 1. Default is 30 seconds.
3. Configure other resources in the service group, if required.
4. Bring the ServiceMonitor resource, and other resources in the service group, online.

Configuring processes

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

Configure a Process resource either manually, as described below, or by using the Application Configuration Wizard.

See "Configuring applications using the Application Configuration Wizard" on page 343 for instructions.

Prerequisites

- The executables configured as the start, stop, and monitor programs must reside on local drives.
- When defining the StartProgram, StopProgram, or MonitorProgram attributes, enclose the path of the executable file in double quotes. *Do not enclose arguments in double quotes*. For example, specify the StartProgram attribute in the following format:

  `StartProgram = "executable_pathname" arguments`

Configuration tasks

1. In your service group, create a resource of type Process.
   See "Adding a resource" on page 200 for instructions.
2. Configure the following required attribute for the resource.
   - StartProgram: The process to be monitored by the agent. You must specify the complete path of the executable, its file extension, and command-line arguments, if any. If you define the start program as a script to launch another program, you must specify the monitor program in the configuration file.
     If you define the start program as a script (a perl script, or a vbs script), the start program should be the program that interprets the script (perl.exe, or cscript.exe) and the script itself should be passed as an argument.

   Configure the following optional attributes, if required:

   - StartupDirectory: The startup directory for the process indicated by the StartProgram attribute.
   - MonitorProgram: A program that monitors the process specified as the start program. You must specify the complete path of the executable, its file extension, and command-line arguments, if any. If you do not

specify a value for this attribute, VCS monitors the start program. However, if the start program is a script to launch another program, you must specify a monitor program.

- MonitorProgramTimeout: The maximum wait time, in seconds, for the agent to receive a return value from the monitor routine. This attribute is ignored if the monitor program is not specified.

3 Configure other resources in the service group, if required.

4 Bring the Process resource, and other resources in the service group, online.

Configuring the MSMQ resource

The following section describes how to configure an MSMQ resource. Before you proceed, review the resource types and attribute definitions of the MSMQ agent in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Prerequisites

- Lanman resource. If MSMQ is integrated with Active Directory, then for the Lanman resource, set the value of the attributes ADUpdateRequired and ADCriticialForOnline to 1.

  Note: You may receive an error when you try to read messages from a remote public queue in Microsoft Message Queuing. See article 889860 in the Microsoft Knowledge Base for more information. To overcome this problem, set the value of the attributes DNSUpdateRequired and DNSCriticialForOnline to 1, for the Lanman resource.

- IP resource

  Note: Ensure that there is only one IP resource per MSMQ resource. If there are multiple MSMQ resources that use the same IP resource, only one MSMQ resource will remain online, and the rest will go into an unknown state.

- Verify that all the existing services that are dependent on the default MSMQ service are in the STOPPED state.
- If MSMQ is installed in Domain Mode, perform the following steps before you bring the MSMQ resource online for the first time:
  - First, bring the Lanman resource online in the Service Group.
  - Next, in the Active Directory, enable the ‘Create All Child Objects’ privilege for the VCS Helper Service user account (HAD Helper) on the MSMQ virtual server.

    Note: You do not need to add this privilege if the VCS Helper Service user account belongs to the Domain Admins group.

To add an MSMQ resource

1. Start VCS on all systems.
2. Ensure that all the required resources are online.
3. Run the MSMQ configuration utility for VCS.

At the command prompt type:

```
msmqconfig -c -n <msmq resource name> -s <nodes in service group
system list> -m <storage path>
```

For example, if the MSMQ resource name is vxmsmq, nodes in the system list are S1 and S2, and the storage path is X:\MSMQ\Storage, then the command will be as follows:

```
msmqconfig -c -n vxmsmq -s S1 S2 -m X:\MSMQ\Storage
```

This will accomplish the preconfiguration and create the registry key HKLM\Software\Microsoft\MSMQ\Clustered QMs\MSMQ$vxmsmq for RegRep.

The MSMQ configuration utility stops and restarts the default MSMQ service.

4 From the Java Console, add a MountV resource (in case of SFW HA) or a NetAppSnapDrive resource (in case of VCS for NetApp SnapMirror) for X:.

   A VMDg resource (in case of SFW HA) or a NetAppFiler resource (in case of VCS for NetApp SnapMirror) may be required if the existing VMDg or NetAppFiler resources do not have a spare volume or LUN to host X:.

5 From the Java Console add a RegRep resource with the registry key created in step 3.

   Add a MountV and a VMDg resource (in case of SFW HA) or a NetAppSnapDrive and a NetAppFiler resource (in case of VCS for NetApp SnapMirror), if they do not exist already, for storing the registry replication information required by RegRep.

   See the *Veritas Cluster Server Bundled Agents Guide* for more information.

6 From the Java Console, add an MSMQ resource and set the following:

   - IPResName to the existing IP resource name
   - LanmanResName to the existing Lanman resource name
   - StoragePath to the storage path for MSMQ. (Referring to the example in step 3, the storage path will be \\MSMQ\\Storage)
   - MountResName to the MountV resource (in case of SFW HA) or FilerResName to the NetAppSnapDrive resource (in case of VCS for NetApp SnapMirror)that you added in step 4.

   You must disable and enable the MSMQ resource every time you make changes to the IPResName attribute.

7 Bring the MSMQ service group online.

8 Run the virtual MMC Viewer on the node where the MSMQ resource is online.

   At the command prompt, type:

```
C:\>VCSVMMCView.exe -target MSMQ
```

You will be prompted for the VirtualName of the LanmanResName provided in the MSMQ resource.

9 Click **Yes** to launch the virtual server MMC.
You can now create, delete, and modify message queues on the virtual MSMQ.
Cluster support for MSMQ triggers is not available in this release. In domain mode MSMQ installation (MSMQ 3.0), if Routing Support is selected while installing MSMQ, it is not supported.

Configuring infrastructure and support agents

This section provides an overview of the steps in configuring the VCS infrastructure and support agents.

- To configure notification in your cluster, use the NotifierManager (NotifierMngr) agent.
  See "Configuring notification" on page 335 for instructions.
- To configure registry replication, use the Registry Replication (RegRep) agent.
  See "Configuring registry replication" on page 336 for instructions.
- To configure web applications, use the VRTSWebApp agent.
  See "Configuring Veritas Web applications" on page 339 for instructions.
- To configure a dummy resource to mirror the state of another resource, use the Proxy agent.
  See "Configuring a proxy resource" on page 340 for instructions.
- To determine the state of a service group with no *OnOff* resources, use the Phantom agent.
  See "Configuring a phantom resource" on page 340 for instructions.
- To test VCS functionality, configure the various file agents.
  See "Configuring file resources" on page 341 for instructions.
- To monitor or manage a service group that exists in a remote cluster, use the RemoteGroup agent.
  See "Configuring a RemoteGroup resource" on page 341 for instructions.

Configuring notification

Use the NotifierMngr agent to set up notification in your cluster. For more information about how VCS handled notification, see "VCS event notification" on page 511. VCS provides a wizard to set up notification, described in "Setting up VCS event notification using the Notifier wizard" on page 227.

Configuring registry replication

The Registry Replication (RegRep) agent replicates the registry of the active cluster node.

Configuration tasks

1. Configure an exclusive MountV resource (in case of SFW HA) or a NetAppSnapDrive resource (in case of VCS for NetApp SnapMirror) for the Registry Replication agent. Verify that no other applications use this resource.
   See "Configuring shared storage" on page 278 for instructions.
2. Create a resource of type RegRep.
   See "Adding a resource" on page 200 for instructions.
3. Configure the following required attributes for the resource.
   - **Keys:** The list of registry keys to be monitored. From the 'name-value' pair of a registry key, you must provide the name of the registry keys to be synchronized and not the value for that key.
     When defining the keys, you must use the abbreviations listed in "Configuring registry keys" on page 337. For excluding certain keys from being replicated, see the instructions listed in "Excluding keys" on page 337. For instructions on replicating registry keys without replicating the subkeys, see "Ignoring subkeys" on page 338.

   Do not configure more than 63 keys for a single RegRep resource otherwise the resource will go in an UNKNOWN state.
   - **MountResName or FilerResName:** The name of the MountV resource (in case of SFW HA) or NetAppSnapDrive resource (in case of VCS for NetApp SnapMirror) on which the Registry Replication agent depends. The resource specifies the mount drive or LUN on the shared disk where the log file is created.
   - **ReplicationDirectory:** The directory on the shared disk in which the registry changes are logged.
4. Configure other resources for the service group, if required.
5. Link the RegRep and MountV (in case of SFW HA) or NetAppSnapDrive (in case of VCS for NetApp SnapMirror) resources such that the RegRep resource depends on the MountV (in case of SFW HA) or NetAppSnapDrive (in case of VCS for NetApp SnapMirror) resource.
   See "Linking resources" on page 213 for instructions.
6. Bring the RegRep resource, and other resources in the service group, online.

Configuring registry keys

To configure a registry key to be replicated or excluded, use the abbreviation corresponding to the registry hive, as listed below:

Table 9-6 RegRep agent - Registry hive and abbreviations

Registry Hive	Abbreviation
HKEY_LOCAL_MACHINE	HKLM
HKEY_CURRENT_USER	HKCU
HKEY_USERS	HKU
HKEY_CURRENT_CONFIG	HKCC
HKEY_CLASSES_ROOT	HKCR

Excluding keys

This section describes the algorithm the Registry Replication agent uses while excluding keys. For example, assume a registry key KEY_X has a subkey of KEY_Y, which has a subkey KEY_Z. This key would appear as KEY_X\KEY_Y\KEY_Z in the Registry Editor. The following table describes various scenarios of keys marked for replication and for exclusion. The Result column describes the agent behavior in these scenarios.

Table 9-7 RegRep agent - Exclude keys and behavior

Keys for Replication	Exclude Keys	Result
KEY_X	KEY_Y\KEY_Z	KEY_Y is excluded. So is KEY_Z.
KEY_X	KEY_Y	KEY_Y is excluded. So is KEY_Z.
KEY_X	KEY_X	KEY_X is *not* excluded and an error message is logged.
KEY_X\KEY_Y	KEY_X	KEY_X is *not* excluded and an error message is logged.

Ignoring subkeys

Use the *IgnoreSubKeys* option for the *Keys* attribute to prevent the RegistryReplication agent from replicating the subkeys. The following table describes possible combination of values for the Keys attribute. The Result column describes the agent behavior in these scenarios:

Table 9-8 RegRep agent - IgnoreSubKeys and behavior

Value Specified for “Keys” Attribute	Result
"HKLM\SOFTWARE\VERITAS\VCS"	Replicates the subkeys
"HKLM\SOFTWARE\VERITAS\VCS"=IgnoreSubKeys	Does not replicate the subkeys
"HKLM\SOFTWARE\VERITAS\VCS"=IgnoreSubKeys:Yes	Does not replicate the subkeys
"HKLM\SOFTWARE\VERITAS\VCS"=IgnoreSubKeys:No	Replicates the subkeys
"HKLM\SOFTWARE\VERITAS\VCS"=<any other value>	Replicates the subkeys

Additional considerations for using IgnoreSubKeys

Symantec recommends not to set the *IgnoreSubKeys* value when the RegRep resource is online. Even if the value is set with the resource online, the changes will be applicable after the next online routine.

Configuring Veritas Web applications

The VRTSWebApp agent brings web applications online, takes them offline, and monitors their status. The agent provides high availability to Veritas Web applications, which support Web consoles of various products.

Configuration tasks

1. In your service group, create resources of type NIC and IP.
   See “Configuring IP addresses in single-adapter systems” on page 283 for instructions.
2. Create a resource of type VRTSWebApp.
3. Configure the following required attributes for the resource.
   - AppName: Name of the application as it appears in the Web server. For example, for VCS, use `vcs`.
   - InstallDir: Path to the Web application installation. The Web application must be installed as a .war file with the same name as the AppName parameter. This attribute should point to the directory that contains this .war file.
   - TimeForOnline: The time the Web application takes to start after it is loaded into the Web server. This parameter is returned as the exit value of the online script, which inform VCS of the time it needs to wait before calling monitor routine on the Web application resource. Typically, this attribute is at least five seconds.
4. Configure other resources for the service group, if required.
5. Link resources to create the following dependencies:
   - IP resource depends on the NIC resource.
   - VRTSWebApp resource depends on the IP resource.

   See “Linking resources” on page 213 for instructions.
6. Bring the VRTSWebApp resource, and other resources in the service group, online.

Configuring a proxy resource

The Proxy agent monitors and mirrors the state of a resource on a local or remote system in a VCS cluster. Use this agent to reduce overheads in configurations where multiple resources point at the same physical device. For example, if multiple service groups use the same NIC, configure one service group to monitor the NIC and have Proxy resources in the other service groups to mirror the state of the NIC resource.

Configuration tasks

1. Create a resource of type Proxy.
   See "Adding a resource" on page 200 for instructions.
2. Configure the following required attribute for the resource:
   - TargetResName: The name of the target resource whose status is to be monitored and mirrored by the Proxy resource.

   If required, configure the following optional attribute for the resource:
   - TargetSysName: The name of the system associated with the target resource. If this attribute is not specified, the Proxy resource assumes the system is local.
3. Configure other resources for the service group, if required.
4. Bring the Proxy resource, and other resources in the service group, online.

Configuring a phantom resource

A Phantom resource enables VCS to determine the status of service groups that do not include *OnOff* resources.

Configuration tasks

1. Create a resource of type Phantom.
   See "Adding a resource" on page 200 for instructions.
2. Configure other resources for the service group, if required.
3. Bring the Phantom resource, and other resources in the service group, online.

Configuring file resources

The FileNone, ElifNone, FileOnOff, and FileOnOnly agents help you test VCS functionality. The process of configuring these resources is similar.

- The FileNone agent monitors a file and returns ONLINE if the file exists.
- The ElifNone agent monitors a file and returns ONLINE if the file does not exist.
- The FileOnOff agent creates, removes, and monitors a file.
- The FileOnOnly agent creates and monitors a file.

Configuration tasks

1. In your service group, create a resource of the desired type.
   See “Adding a resource” on page 200 for instructions.
2. Configure the following required attribute for the resource.
   - PathName: The complete path of the file to be monitored.
3. If required, configure additional resources in the service group.
4. Bring the file resource, and other resources, in the service group online.

Configuring a RemoteGroup resource

The RemoteGroup agent establishes dependencies between applications that are configured on different VCS clusters. With the RemoteGroup agent, you can monitor or manage a service group that exists in a remote cluster. Some points about configuring the RemoteGroup resource are:

- For each remote service group that you want to monitor or manage, you must configure a corresponding RemoteGroup resource in the local cluster.
- Multiple RemoteGroup resources in a local cluster can manage corresponding multiple remote service groups in different remote clusters.
- You can include the RemoteGroup resource in any kind of resource or service group dependency tree.
- A combination of the state of the local service group and the state of the remote service group determines the state of the RemoteGroup resource.
- Global groups are not supported as remote service groups.

Before configuring the RemoteGroup resource, review the resource types, the attribute definitions, and the sample scenario described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Configuration tasks

1 In your service group, create resources of type IP and NIC.
See "Adding a resource" on page 200 for instructions.

2 Create a resource of type RemoteGroup.
See "Adding a RemoteGroup resource from the Java Console" on page 203 for instructions.

3 Configure the required attributes for the RemoteGroup resource. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on the required attributes and their definitions.

4 Link the resources as follows:
   - Link the IP and NIC resources such that the IP resource depends on the the NIC resource.
   - Link the RemoteGroup and NIC resources such that the RemoteGroup resource depends on the NIC resource.

   See "Linking resources" on page 213 for instructions.

5 Configure other resources in the service group, if required.

6 Bring the IP, NIC, and RemoteGroup resources online.

Configuring applications using the Application Configuration Wizard

Veritas Cluster Server provides an Application Configuration Wizard to create service groups to monitor applications, which are configured as resources of type GenericService, ServiceMonitor, or Process. The wizard also enables you to add file share, registry replication, and network resources to application service groups.

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

Prerequisites

- Verify that the binaries of the application to be configured are present on the nodes on which the service group will be configured.
- Verify that the shared drives or LUNs required by the applications are mounted.
- Before running the wizard, make sure you have the following information ready:
  - Type of applications for which resources are to be configured.
  - Shared storage used by the applications.
  - Registry replication information.
  - Network information.

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

Adding resources to a service group

1 Start the Application Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Configuration Wizard**)

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

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

4 Specify the service group name and system list.

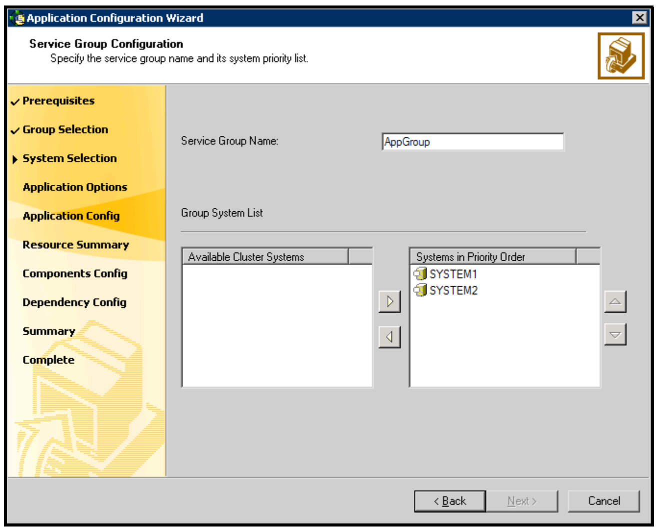

- Enter a name for the service group.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list.
  To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
- To change a system's priority in the service group's system list, select the system in the Systems in Priority Order list and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

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

- Generic Service: Configures a service using the Generic Service agent. The agent brings services online, takes them offline, and monitors their status.
  See "Configuring a GenericService resource" on page 345.
- Process: Configures a process using the Process agent. The agent brings processes online, takes them offline, and monitors their status.
  See "Configuring processes" on page 349.
- Service Monitor: Configures a service using the ServiceMonitor agent. The agent monitors a service or starts a user-defined script and interprets the exit code of the script.
  See "Configuring a ServiceMonitor resource" on page 353.

Configuring a GenericService resource

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

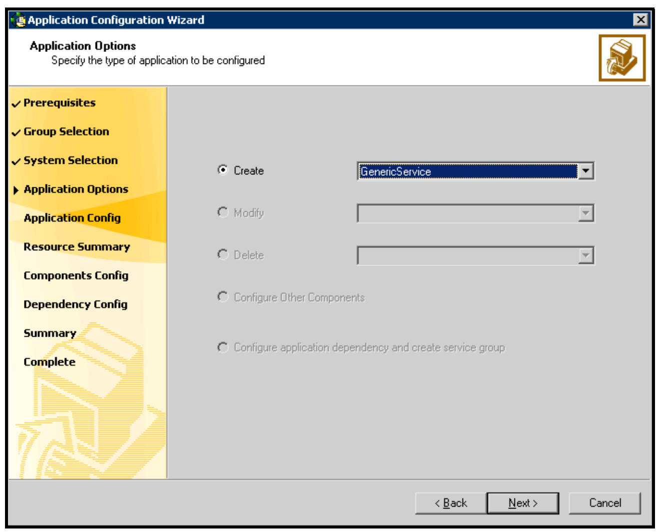

2 Select the service name for which you wish to configure a Generic Service resource. Also specify the attributes for the resource.

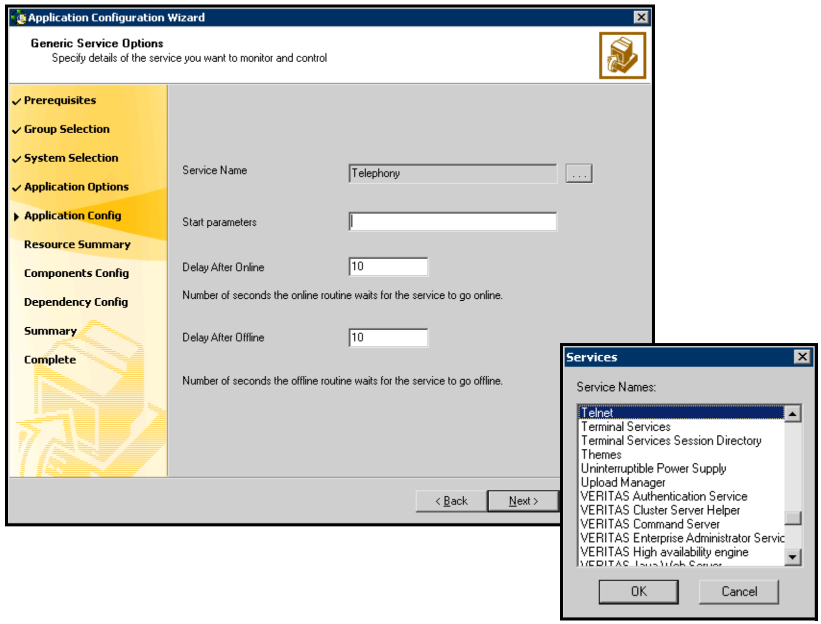

- Click the icon (...) adjacent to the Service Name text box.
- In the Services dialog box, select a service and click **OK**. The selected service appears in the Service Name text box.
- In the Start Parameters text box, provide the start parameters for the service, if any.
- In the Delay After Online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

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

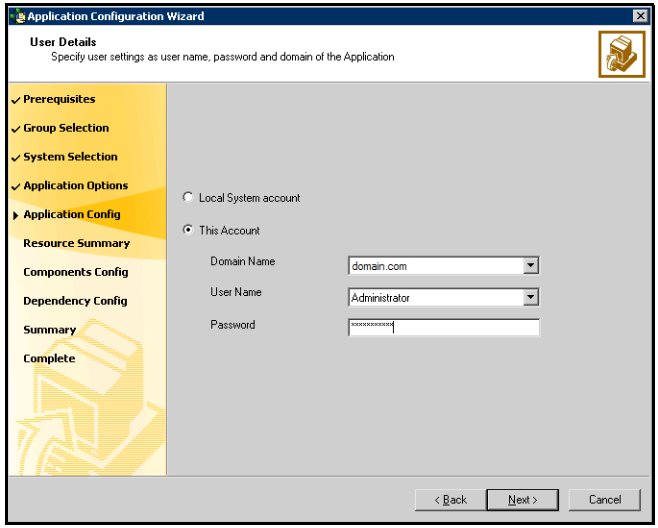

- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

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

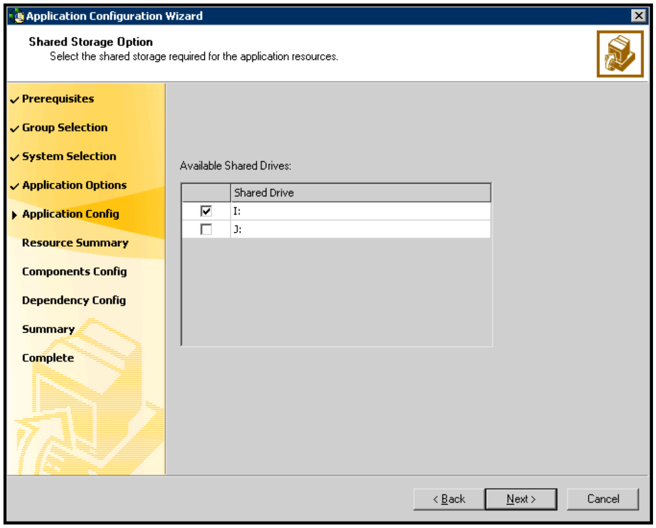

- In the Available Shared Drives box, select the check box adjacent to the shared drive.
- Click **Next**.

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

6 The Application Options dialog box appears. Select one of the following options:

- To configure another GenericService resource, repeat step 1 through step 5.
- To configure a Process resource, proceed to "Configuring processes" on page 349 for instructions.
- To configure a ServiceMonitor resource, proceed to "Configuring a ServiceMonitor resource" on page 353 for instructions.
- To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to "Configuring VCS components" on page 356 for instructions.

If you do not want to add any more resources to your service group, proceed to "Creating service groups" on page 362.

Configuring processes

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

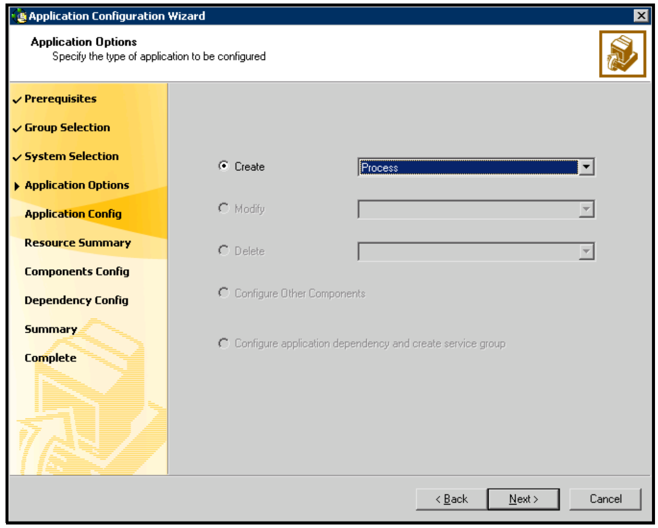

2 Specify the details for the process.

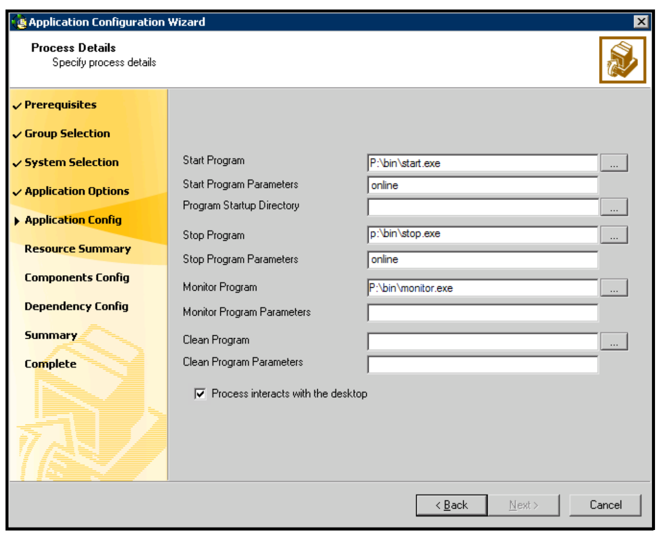

- In the Start Program text box, specify the complete path of the program that will start the process to be monitored by VCS. You can choose to either type in the location of the program or browse for it using the (...) icon.

- In the Start Program Parameters text box, specify the parameters used by the Process agent start program.
- In the Program Startup Directory text box, type the complete path of the Process agent program or browse for it by clicking the (...) icon.
- In the Stop Program text box, type the complete path of the program that will stop the process started by the Start Program or browse for it by clicking the (...) icon.
- In the Stop Program Parameters text box, specify the parameters used by the stop program.
- In the Monitor Program text box, type the complete path of the program that monitors the Start Program or browse for it by clicking the (...) icon.
  If you do not specify a value for this attribute, VCS monitors the Start Program. If the Start Program is a script to launch another program, you must specify a monitor program.
- In the Monitor Program Parameters text box, specify the parameters used by the monitor program.
- In the Clean Program text box, type the complete path of the Clean process or browse for it by clicking the (..) icon.
  If no value is specified, the agent kills the process indicated by the Start Program.
- In the Clean Program Parameters text box, specify the parameters used by the Clean program.
- Check the **Process interacts with the desktop** check box if you want the process to interact with your Windows desktop. Setting this option enables user intervention for the process.
- Click **Next**.

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

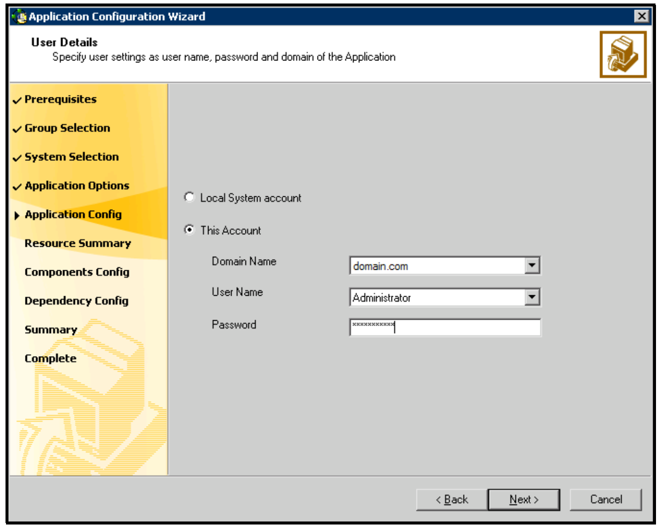

- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

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

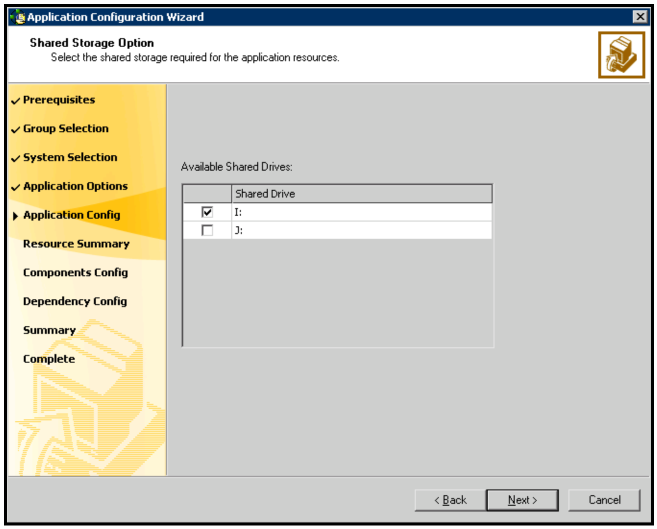

- From the Available Shared Drives box, select the check box adjacent to the shared drive.
- Click **Next**.

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

6 The Application Options dialog box appears. Select one of the following options:

- To configure another Process resource, repeat step 1 through step 5.
- To configure a GenericService resource, see "Configuring a GenericService resource" on page 345 for instructions.
- To configure a ServiceMonitor resource, proceed to "Configuring a ServiceMonitor resource" on page 353 for instructions.
- To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to "Configuring VCS components" on page 356 for instructions.

If you do not want to add any more resources to your service group, proceed to "Creating service groups" on page 362.

Configuring a ServiceMonitor resource

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

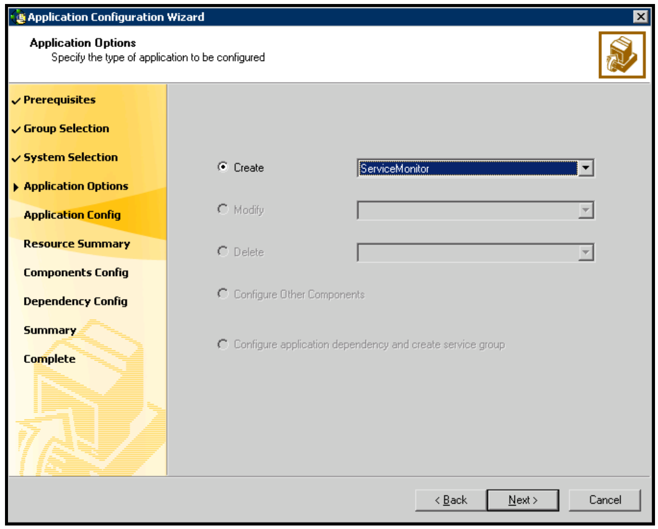

2 Specify the service to be monitored or a user-defined script to monitor a service.

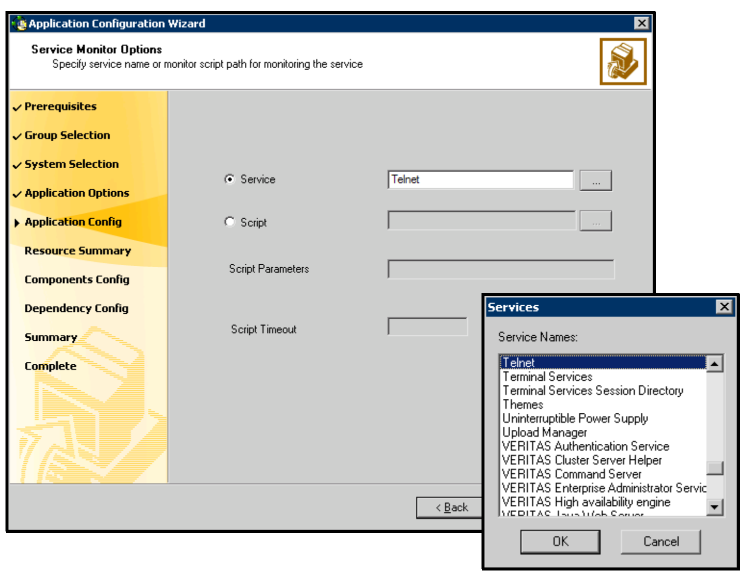

If you want VCS to monitor the service:

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

If you want a script to monitor the service:

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

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

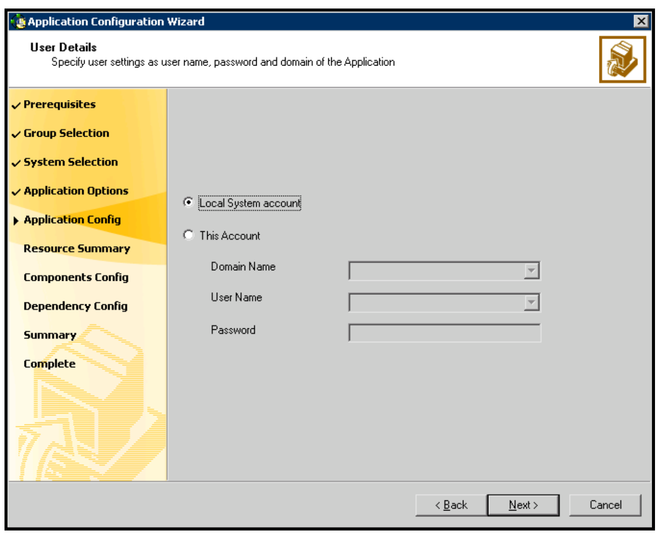

- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
  If the service selected in step 2 on page 353 is running in the context of a local system account, the **This Account** option is disabled. Similarly, if

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

- Click **Next**.

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

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

5 The Application Options dialog box appears. Select one of the following options:

- To configure another ServiceMonitor resource, repeat step 1 through step 4.
- To configure a GenericService resource, see “Configuring a GenericService resource” on page 345 for instructions.
- To configure a Process resource, see “Configuring processes” on page 349 for instructions.
- To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “Configuring VCS components” on page 356 for instructions.

If you do not want to add any more resources to your service group, proceed to “Creating service groups” on page 362.

Configuring VCS components

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

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

To configure VCS components

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

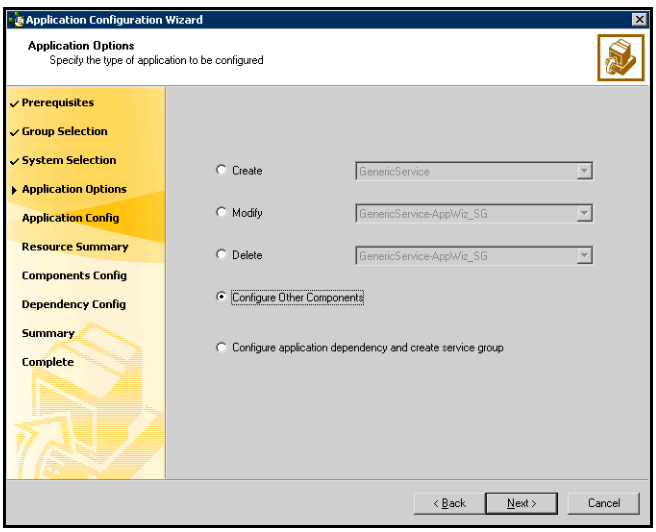

2 Select the VCS component to be configured for your applications.

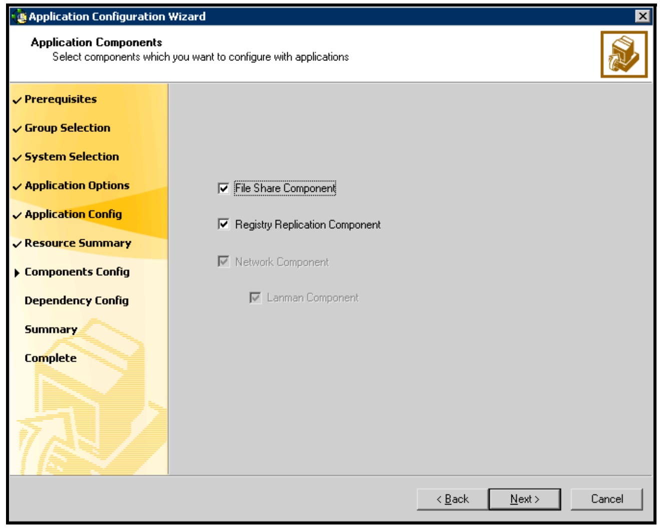

The available options are:

- File Share Component: Select this to configure a FileShare resource for your application. To configure a FileShare resource, proceed to the next step.
  If you select to configure the file share component, the **Network Component** check box is checked by default.
- Registry Replication Component: Select this option to configure registry replication for your application. To configure a Registry Replication resource, proceed to step 5 on page 360.
- Network Component: Select this option to configure network components for your application. If you wish to configure a virtual computer name, check **Lanman component** also. To configure a network resource, proceed to step 7 on page 361.

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

To configure a FileShare resource

3 Specify the configuration information for the FileShare resource to be created.

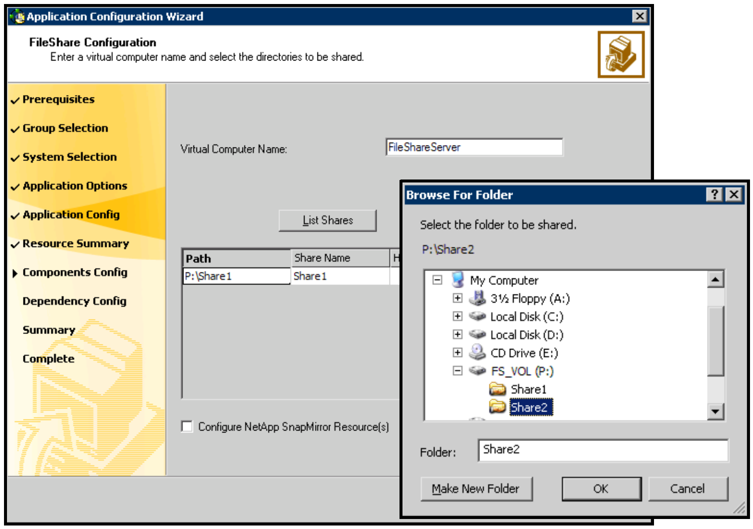

The FileShare Configuration dialog box appears only if you chose FileShare component in the Application Component dialog box.

- Enter a unique virtual computer name by which the server will be known to clients. The virtual name must not exceed 15 characters.
- Click **List Shares** to view all the existing shares on the shared storage. Select a share and click **Add**.
- In the **Path** column, type or click the path of the directories to be shared. Click the Edit icon (**...**) to browse for folders. The selected directories must meet the following conditions:
  - The selected drive, the mount path, and the file path must not exist in the VCS configuration.
  - The directories to be shared must reside on shared, non-system drives.

  The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.
- If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can selected a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.

- To make the new share a hidden share, check **Hide Share**. To share subdirectories, check **Share Subdirs**. To hide shared subdirectories, check **Hide Child Shares**.
- Click **Add** to add a file share.

Repeat these steps for each file share to be created. Click **Remove** to remove a file share from the configuration.

- Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. This is applicable in case of VCS for NetApp SnapMirror only. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.
- Click **Next**. The wizard begins validating your configuration. Various messages indicate the validation status. After the validations are completed, the Share Permissions dialog box appears.

4 Specify the users for the file share and assign permissions to them.

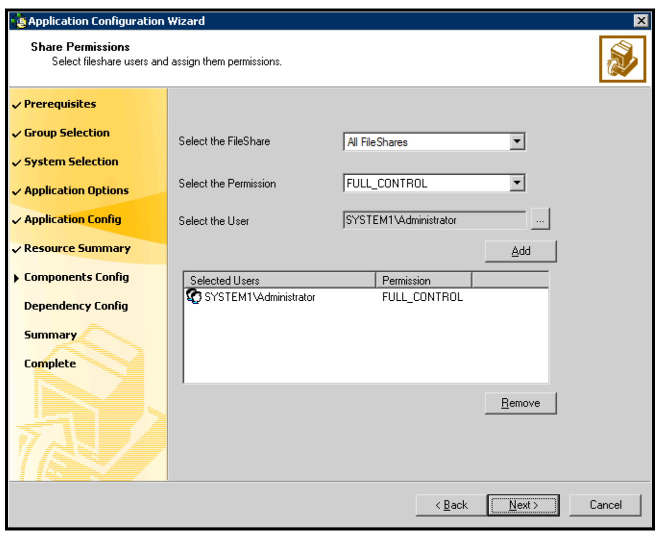

- In the Select the FileShare list, click the file share with which to associate user permissions or click the default **All FileShares** to set the same permissions for all file shares.
- In the Select the Permission list, click the permission to be associated with the user.

- In the Select the User field, click the **..** button and specify the users to be assigned permissions for the selected file share. Users not selected here will not have access to the file share.
- Click **Add**. The selected users are transferred to the Selected Users list. By default, all selected users are given READ_ACCESS permission.
- In the Selected Users list, click the user whose default permission you want to change. In the Select the Permission list, click the permission to be associated with the user. To deny file share access to a user, click the user name in the Selected Users list and click **Remove**.

  Repeat the process for each file share with which to associate users and permissions.
- Click **Next**.

If you selected Registry Replication from the Application Component dialog box, proceed to the next step. Otherwise, proceed to step 7 on page 361.

To configure Registry Replication

5 Specify the registry keys to be replicated.

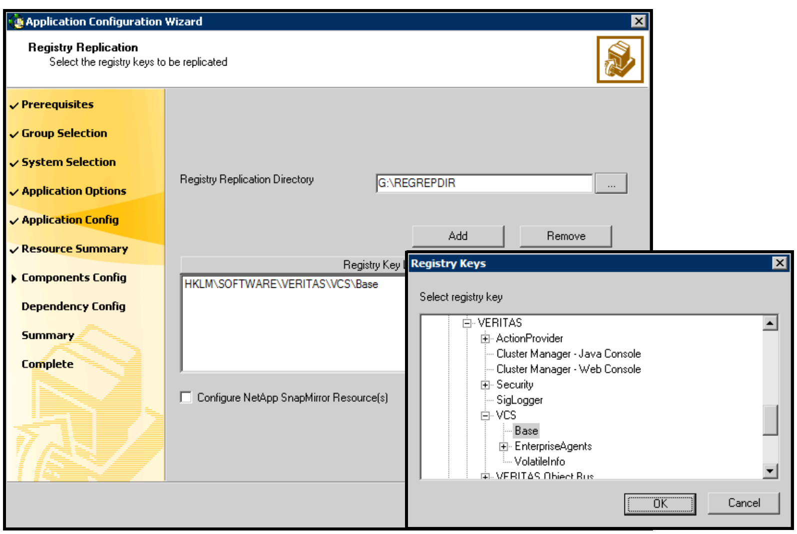

The RegistryReplication dialog box appears only if you chose to configure the Registry Replication Component in the Application Component dialog box.

- Specify the directory on the shared disk in which the registry changes are logged.

- Click **Add**.
- In the Registry Keys dialog box, select the registry key to be replicated.
- Click **OK**. The selected registry key is added to Registry KeyList box.
- Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. This is applicable in case of VCS for NetApp SnapMirror only. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.
- Click **Next**.

If you chose Network Component from the Application Component dialog box, proceed to the next step. Otherwise, proceed to step 8 on page 362.

To configure network components

6 This is applicable in case of VCS for NetApp SnapMirror only.

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

If you are configuring MPIO over FC, you must select at least 2 FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

7 Specify information related to your network.

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

- Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 15 characters.

  Note that the Virtual Computer Name text box is displayed only if you chose to configure the Lanman Component in Application Component dialog box. However, if you chose to configure a FileShare resource in the service group, the Virtual Computer Name text box is not displayed. In such a case, the Lanman resource uses the virtual computer name specified for the FileShare resource.
- Enter a unique virtual IP address for the virtual server.
- Enter the subnet to which the virtual server belongs.
- Click **Advanced...** to specify additional details for the Lanman resource.

- Check **AD Update required** to enable the Lanman resource to update the Active Directory with the virtual name.
- Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

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

- Click **OK**.

- For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.
  Note that the wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private.
- Click **Next**.

8 The Application Options dialog box is displayed. Select one of the following options:
- To configure additional VCS components, repeat step 1 on page 356 through step 7 on page 361.
- To configure a GenericService resource, see “Configuring a GenericService resource” on page 345 for instructions.
- To configure a Process resource, see “Configuring processes” on page 349 for instructions.
- To configure a Service Monitor resource, see “Configuring a ServiceMonitor resource” on page 353 for instructions.

If you do not want to add any more resources to your service group, proceed to “Creating service groups” on page 362.

Creating service groups

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

To create a service group

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

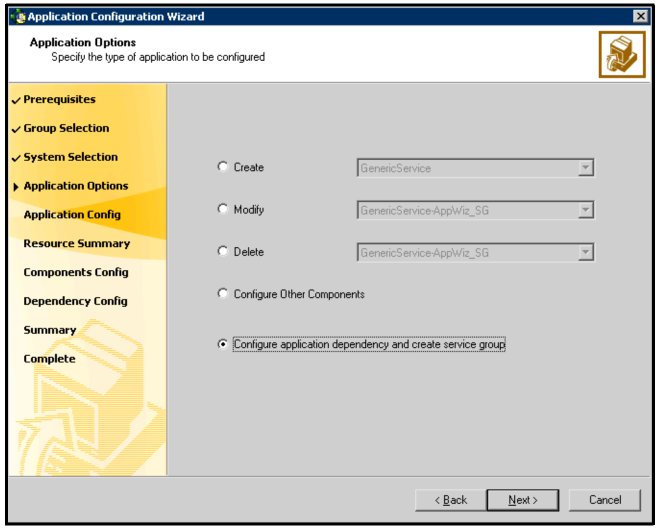

The option is enabled only if:

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

2 Specify the dependency between the applications.

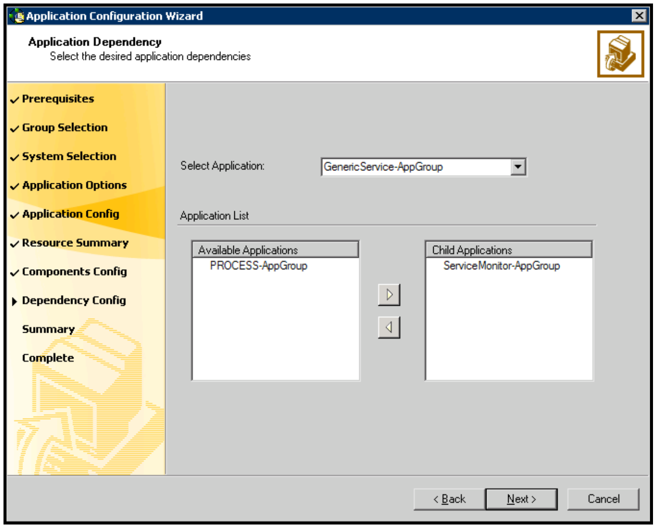

You must have at least two resources configured for this dialog box to appear. Of the two resources, one should either be a GenericService or a Process resource.

- From the Select Application list, select the application that would depend on other applications. The selected application becomes the parent application.
- From the Available Applications list, select the application on which the parent application would depend and click the right-arrow icon to move the application to the Child Applications list.
  To remove an application from the Child Applications list, select the application in the list and click the left arrow.

Repeat these steps for all such applications for which you want to create a dependency.

- Click **Next**.

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

3 Review the service group configuration.

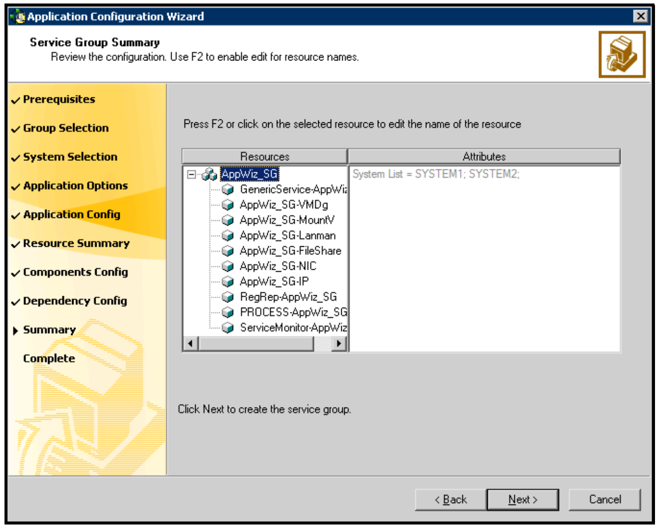

The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

- The wizard assigns unique names to resources. Change names of resource, if required.
  To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.
- Click **Next**.
- A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.
  The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

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

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

Modifying an application service group

This section describes how to modify a service group using the Application Configuration Wizard.

- If the service group to be modified is online, you must run the wizard from a system on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

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

To modify a service group

1. Start the Application Configuration Wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Modify service group**. From the Service Groups list, select the service group containing the resource that you want to modify and click **Next**.
4. In the Service Group Configuration panel, click **Next**.

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

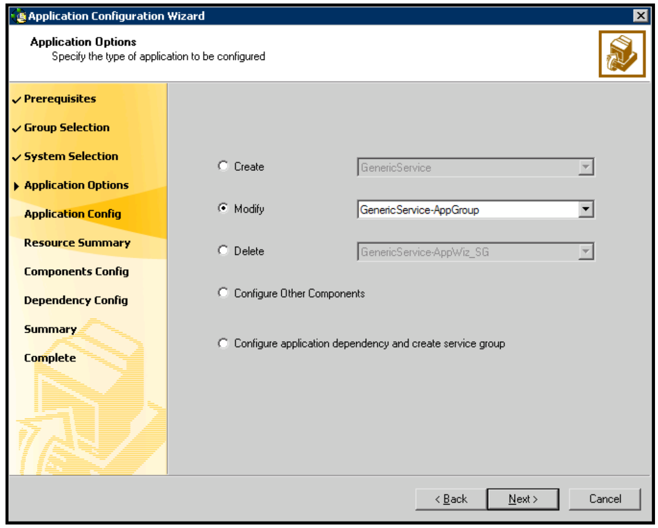

The **Modify** option is enabled only if:

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

6 Depending on the resource you chose to modify from the Application Options page, you would either get the Generic Service Options (see "Configuring a GenericService resource" on page 345), Process Details (see "Configuring processes" on page 349), or Service Monitor Options dialog box (see "Configuring a ServiceMonitor resource" on page 353). Make required changes in the appropriate dialog box and click **Next**.

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

8 In the Application Resource Summary dialog box, review the summary of the resource. Click **Back** to make changes. Otherwise, click **Next**.

9 The Application Options dialog box appears. Repeat step 5 through step 8 for each resource that you want to modify.

10 After modifying the required resources, you can:

- Add additional resources to the service group.
  See "Adding resources to a service group" on page 344 for instructions.
- Delete resources from the service group.
  See "Deleting resources from a service group" on page 369 for instructions.
- Add VCS components to the service group.
  See "Configuring VCS components" on page 356 for instructions.
- Create the service group.
  See "Creating service groups" on page 362 for instructions.

Deleting resources from a service group

This section describes how to delete a resource within a service group using the Application Configuration Wizard.

To delete a resource

1. Start the Application Configuration Wizard. **(Start > All Programs > Symantec Cluster Server > Configuration Wizards > Application Configuration Wizard**)
2. Read the text on the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Modify Service Group**. From the Service Groups list, select the service group containing the resource that you want to delete and click **Next**.
4. In the Service Group Configuration panel, click **Next**.
5. In the Application Options panel, click **Delete**, select the resource you want to delete, and click **Next**.
6. In the Warning dialog box, click **No** to retain the selected resource. Otherwise, click **Yes**.
   The specified resource will be deleted when you exit the wizard after selecting the **Configure application dependency and create service group** option in the Application Options panel.
7. After marking the resource for deletion, you can:
   - Add additional resources to the service group.
     See "Adding resources to a service group" on page 344 for instructions.
   - Modify resources in the service group.
     See "Modifying an application service group" on page 366 for instructions.
   - Add VCS components to the service group.
     See "Configuring VCS components" on page 356 for instructions.
   - Create the service group.
     See "Creating service groups" on page 362 for instructions.

Deleting an application service group

This section describes steps to delete an application service group using the configuration wizard.

To delete a service group

1. Start the Application Configuration Wizard on a system configured to host the application service group. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Application Configuration Wizard**)
2. Review the information in the Welcome panel and click **Next**.
3. In the Wizard Options panel, click **Delete service group**, select the service group to be deleted, and click **Next**.
4. In the Service Group Summary panel, click **Next**.
5. A message appears informing you that the wizard will run commands to delete the service group. Click **Yes** to delete the service group.
6. Click **Finish**.

Modifying the cluster configuration

This chapter describes how to modify and delete a cluster configuration using the VCS Configuration wizard. The chapter also describes how to enable and disable Veritas Security Services in clusters configured to run in secure mode.

Use the VCS Configuration Wizard (VCW) to modify and delete a cluster configuration. When used to modify a cluster configuration, the wizard performs the following tasks:

- Adds nodes to a cluster
- Remove nodes from a cluster
- Reconfigures the private network and LLT
- Reconfigures Veritas Security Services
- Configures the ClusterService service group in the cluster

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

Adding nodes to a cluster

Before adding a node to a cluster, install VCS on the node:

- If you purchased VCS for Network Appliance SnapMirror, refer to the following guides depending on the application:
  - *Veritas Cluster Server Implementation Guide for Microsoft Exchange with NetApp SnapMirror*
    or
    Veritas Cluster Server Implementation Guide for Microsoft Exchange Server 2007 with NetApp SnapMirror
    or
    Veritas Cluster Server Implementation Guide for Microsoft SQL with NetApp SnapMirror
    or
    Veritas Cluster Server Implementation Guide for Oracle with NetApp SnapMirror
- If you purchased Storage Foundation for Windows High Availability, refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

The VCS Configuration wizard configures VCS components and starts VCS services on the new node. The wizard does not configure any service groups on the new node.

Note: To add nodes to single node cluster without private link heartbeat configured, you must first reconfigure the cluster to include the private links. See "Reconfiguring a cluster" on page 384 for instructions.

To add a node to a VCS cluster

1. Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server> Configuration Wizards > 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.

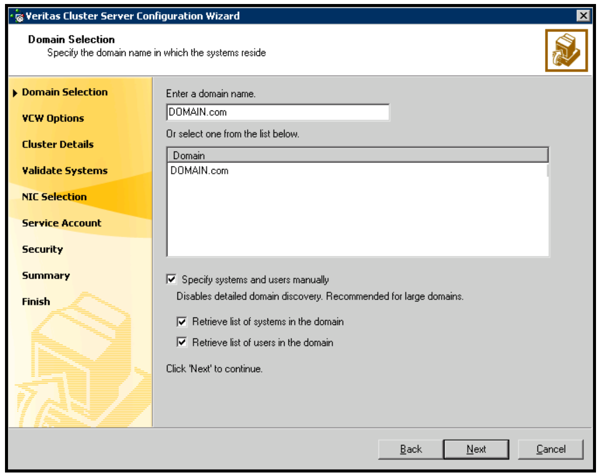

To discover information about all the systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

  Proceed to step 7 on page 376.

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 375. Otherwise proceed to the next step.

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

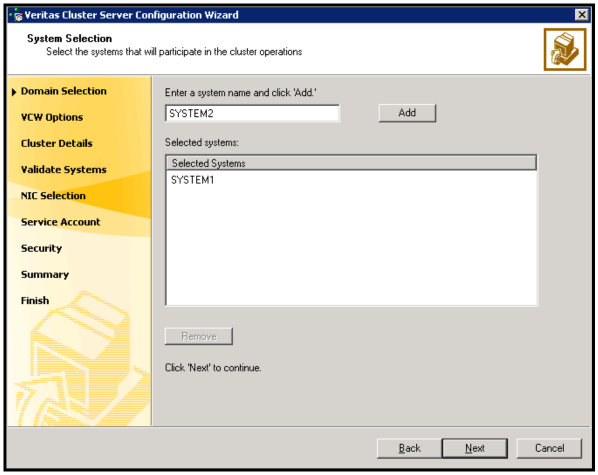

- Type the name of a node in the cluster and click **Add**.
- Type the name of the system to be added to the cluster and click **Add**.

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 7 on page 376.

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.

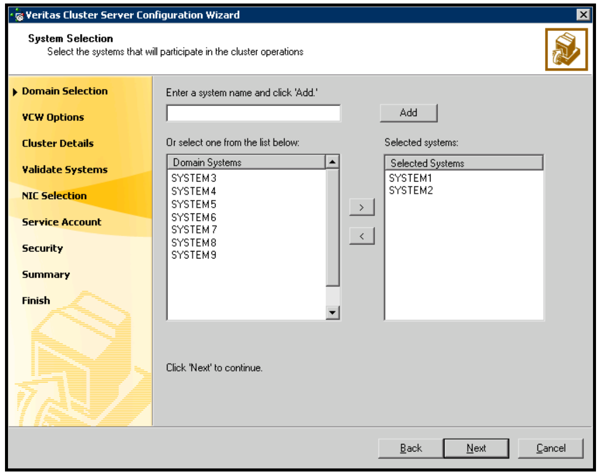

Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

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

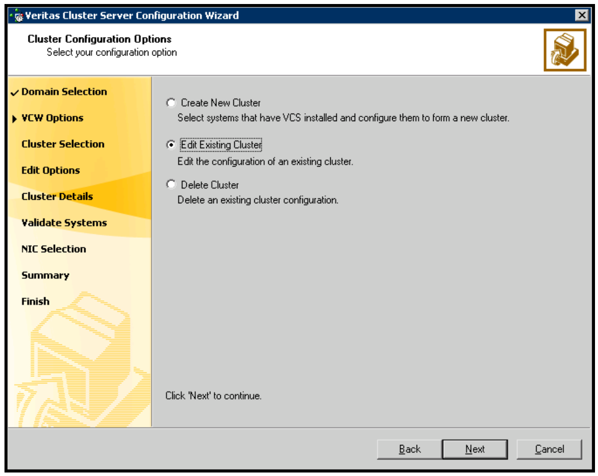

8 On the Cluster Selection panel, select the cluster to be edited and click **Next**.

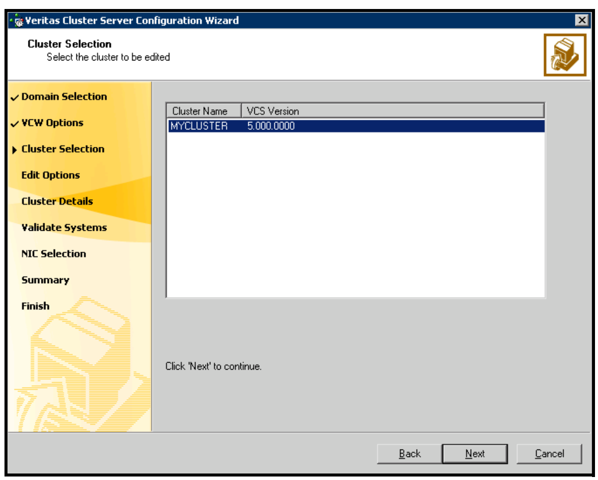

If you chose to specify the systems manually in step 4, only the clusters configured with the specified systems are displayed.

9 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.

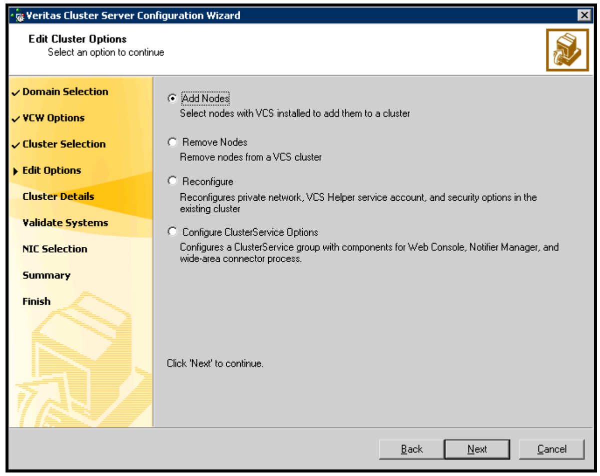

In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

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).

10 On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.

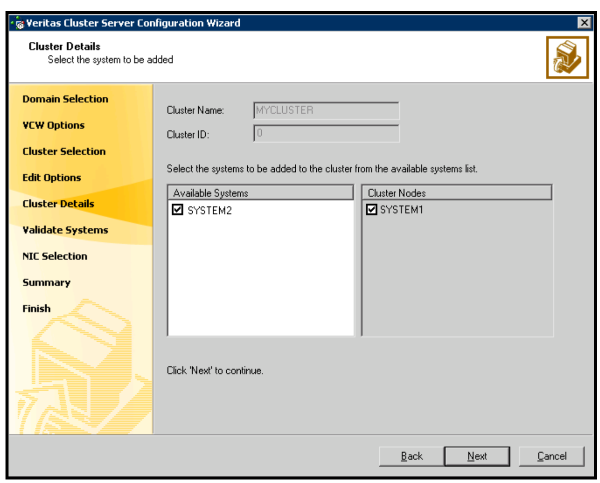

The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

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

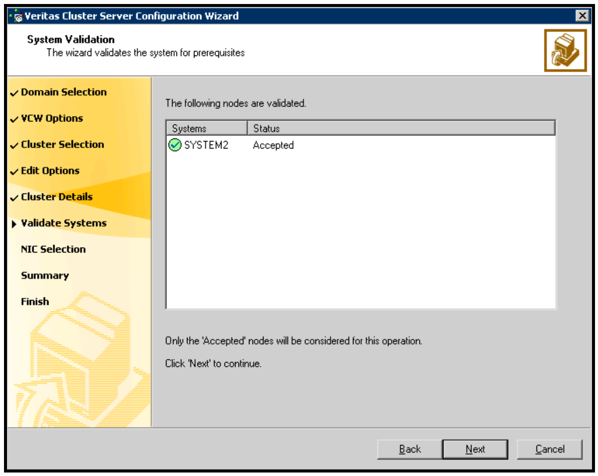

If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

12 On the Private Network Configuration panel, select two NICs for the VCS private network communication, on each system being added, and then click **Next**.

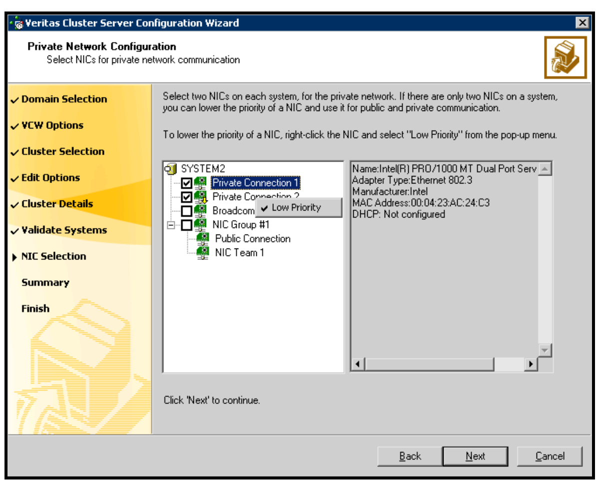

- Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, make sure you lower the priority of the NIC that is used for 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.*

13 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.

14 Specify the credentials for the user in whose context the VCS Helper service runs.

15 Review the summary information and click **Add**.

16 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

Modifying values for ClusterService group attributes

Modify the following ClusterService group attributes on all the newly added nodes to include local values:

- MACAddress attributes of all the NIC resources
- MACAddress attributes of all the IP resources
- StartProgram, StopProgram, and MonitorProgram attributes of the wac resource
- InstallDir attribute of VCSWeb resource

You can modify these values from the VCS Java Console or Web Console.

Removing nodes from a cluster

The following steps describe how to remove nodes from a multiple node VCS cluster. To remove a node from a single node cluster, you must delete the cluster.

See "Deleting a cluster configuration" on page 401 for instructions.

To remove nodes from a cluster

1. Verify no service groups are online on the node to be removed.
2. Remove the node from the SystemList of all service groups.
3. Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
4. Read the information on the Welcome panel and click **Next**.
5. In the Configuration Options panel, click **Cluster Operations** option and click **Next**.
6. In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the domain discovery options.

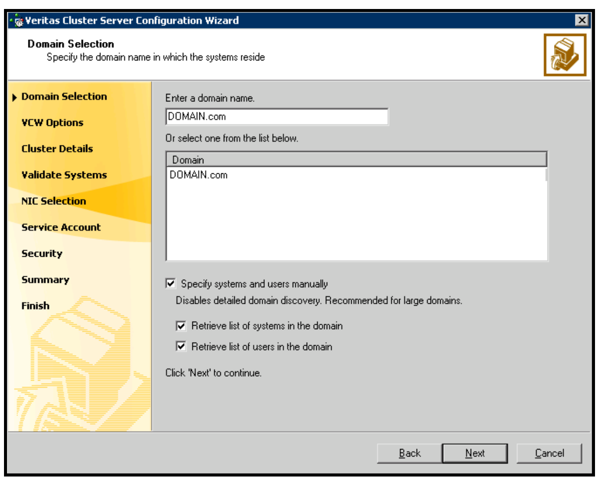

To discover information about all the systems and users in the domain

- Uncheck the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 9 on page 382.

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

- Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
  If you checked **Retrieve system list from domain**, proceed to step 8 on page 382. Otherwise proceed to the next step.

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

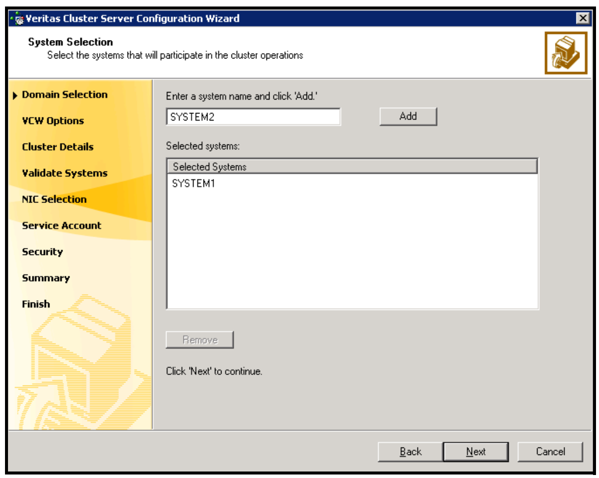

Proceed to step 9 on page 382.

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

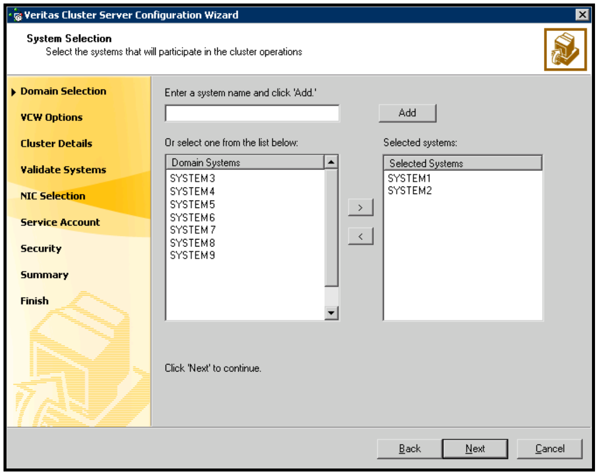

Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster.

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

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

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

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

In the Cluster User Information panel, enter the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you remove a node from a non-secure cluster.

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

The wizard does not remove a root broker node from a secure cluster. To remove a root broker node, you must reconfigure the cluster to have a different root broker.

See “Reconfiguring a cluster” on page 384 for instructions.

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

14 The wizard validates the selected nodes. After the nodes have been validated, click **Next**. If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

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

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

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

Reconfiguring a cluster

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

To reconfigure a cluster

1. Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server> Configuration Wizards > Cluster Configuration Wizard**)
2. Read the information on the Welcome panel and click **Next**.
3. In the Configuration Options panel, click **Cluster Operations** and click **Next**.
4. In the Domain Selection panel, select or type the name of the domain in which the cluster resides and click **Next**.

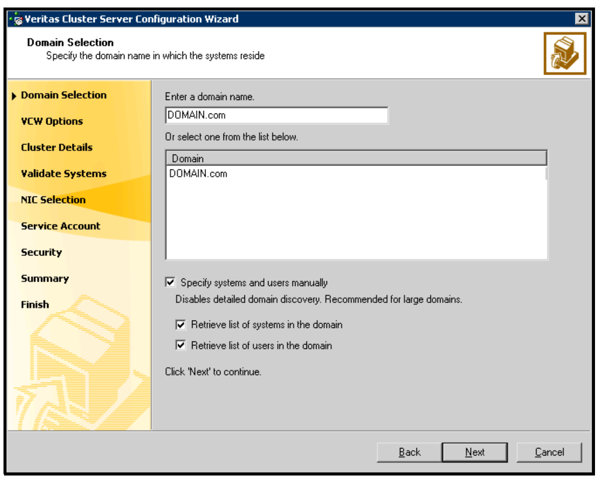

To discover information about all the systems and users in the domain

- Uncheck the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 7 on page 386.

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

- Check the **Specify systems and users manually** check box.
  Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

If you checked **Retrieve system list from domain**, proceed to step 6 on page 385. Otherwise proceed to the next step.

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

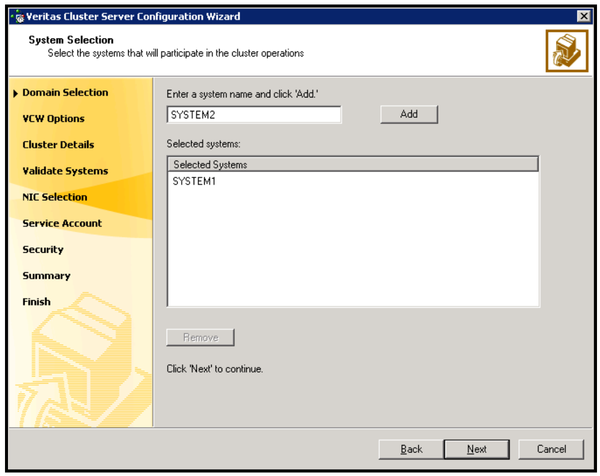

Proceed to step 7 on page 386.

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

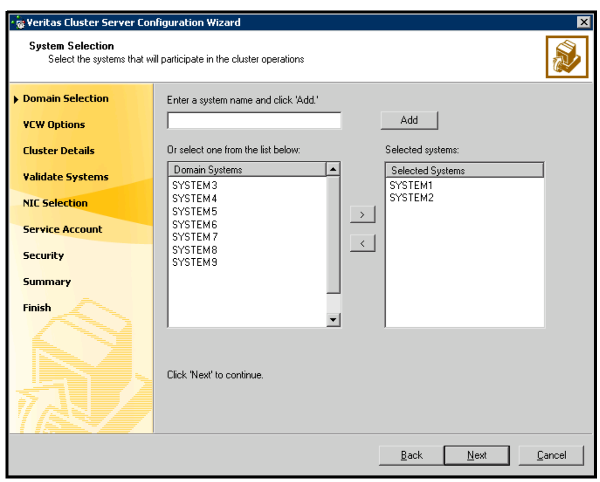

Enter the system name and click **Add** to add the system to the Selected Systems list. Alternatively, you can select the systems from the Domain Systems list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster.

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

8 In the Cluster Selection panel, select the cluster to be reconfigured and click **Next**. If you chose to specify the systems manually in step 4, only the clusters configured with the specified systems are displayed.

9 In the Edit Cluster Options panel, click **Reconfigure** and click **Next**.
In the Cluster User Information dialog box, enter the user name and password for a user with administrative privileges to the cluster and click **OK**.
The Cluster User Information dialog box appears only when you reconfigure a non-secure cluster.

10 In the second Edit Cluster Options dialog box, select any of the following options:

- **Change private network heartbeat links**
  Select this option to change the private network heartbeat links.
  If the selected cluster is a single node cluster, the option is to remove the private heartbeat links.
  If the cluster has more than one node, the options are to add or remove private heartbeat links.
  See step 11.
- **Change HAD Helper User account**
  Selection this options to change the user account for the Veritas Cluster Server Helper service.
  See step 12.
- **Configure/Change Veritas Security Services**
  Select this option to configure or change Veritas Security Services for single sign-on, or to add or change the root broker.
  See step 13.

11 If the option to change the private network heartbeat links was selected, select two NICs on each system for the VCS private network.

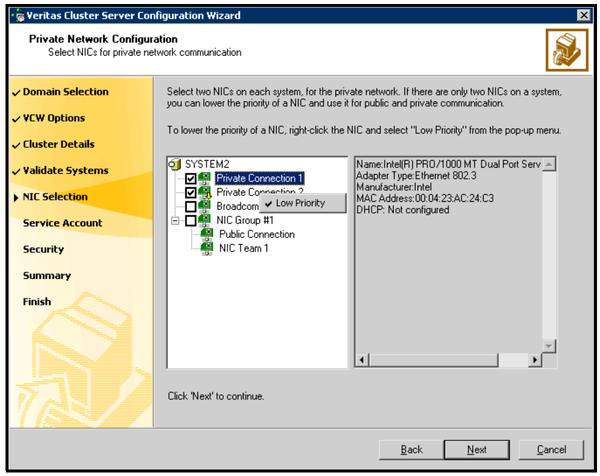

- Select two NICs to be assigned to the private network.
  Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
  If you have only two NICs on a selected system, make sure you lower the priority of at least one NIC for that system. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- Click **Next**.

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.*

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

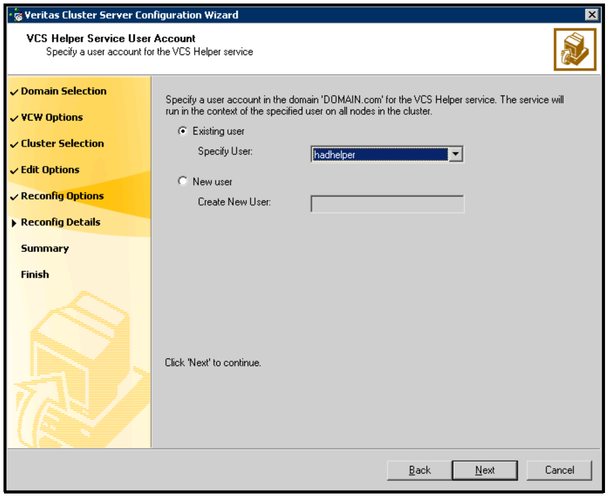

The VCS High Availability Daemon, which runs in the context of the local system built-in account, uses the VCS Helper Service user context to access the network.

- Select one of the following
  - **Existing user**
    Choose an existing user account context for the VCS Helper service.
  - **New user**
    Create a new user account context for the VCS Helper service.
  - **Configured users**
    If multiple user accounts are configured on different cluster nodes, this option appears. Select one of the user accounts to configure. This account will then be configured on all the cluster nodes.
- Enter a valid user name for the selected account and click **Next**.

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

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

13 If the option to change Symantec Product Authentication Service was selected, specify security settings for the cluster.

The wizard enables you to configure Symantec Product Authentication Service. If it is already configured, the wizard enables you to change the root broker.

To configure Symantec Product Authentication Service in the cluster

- Check the **Configure Security Services** check box.
- Enter or select a cluster node or a system in the domain to serve as the root broker.

  If you select a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Note that you cannot make the root broker services highly available.

  If you select a system that is not part of the cluster, make sure the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.
- Click **Next**.

To change the root broker

- Check the **Change Symantec Product Authentication Service Configuration** check box.
- Enter or select a system to be configured as the root broker. A root broker can either be a system within the cluster or a system in the domain. If you select a system that is not a part of the cluster, make sure it is configured to serve as a root broker. All other systems in the cluster are configured as authentication brokers.
- Click **Next**.

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

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

Enabling and disabling Symantec Product Authentication Service

This section describes how to enable and disable Symantec Product Authentication Service for clusters running in secure mode. Note that these procedures apply only if you have configured Symantec Product Authentication Service in the cluster.

For instructions on how to configure Symantec Product Authentication Service in a cluster, see "Reconfiguring a cluster" on page 384.

To enable Symantec Product Authentication Service

1. Stop VCS on all nodes:
   `C:\>` **`hastop -all`**
2. Stop the Veritas Command Server service on all nodes.
3. On each node in the cluster, create an empty file with the name `.secure` under `%VCS_HOME%\conf\config`, where *%VCS_HOME%* represents the VCS installation directory, typically `C:\Program Files\Veritas\Cluster Server`.
4. Start the Veritas Command Server service on all nodes.
5. Set the SecureClus attribute to 1 in your configuration file main.cf.
   Open the configuration file main.cf using Notepad, and add the following line in the cluster definition:
   `SecureClus = 1`
   For example:
   ```
   cluster VCSCluster9495 (
       UserNames = { admin = gmnFmhMjnInnLvnHmk }
       Administrators = { admin }
       SecureClus = 1
       CredRenewFrequency = 0
       CounterInterval = 5
   )
   ```
6. Save the configuration.
7. Start the VCS engine on the local node:
   `C:\>` **hastart**
8. Start VCS on other nodes in the cluster.

To disable Symantec Product Authentication Service

1. Stop VCS on all nodes:
   `C:\> hastop -all`
2. Stop the Veritas Command Server service on all nodes.
3. From each node in the cluster, remove the `.secure` file from `%VCS_HOME%\conf\config`, where *%VCS_HOME%* represents the VCS installation directory, typically `C:\Program Files\Veritas\Cluster Server`.
4. Type the following on the command prompt:
   `C:\> cmdserver -withoutsecurity`
5. Start the Veritas Command Server service on all the nodes, if the service has not started.
6. Set the SecureClus attribute to 0 in your configuration file main.cf.
   Open the configuration file main.cf using Notepad, and add the following line in the cluster definition:
   ```
   SecureClus = 0
   ```
   For example:
   ```
   cluster VCSCluster9495 (
       UserNames = { admin = gmnFmhMjnInnLvnHmk }
       Administrators = { admin }
       SecureClus = 0
       CredRenewFrequency = 0
       CounterInterval = 5
   )
   ```
7. Save the configuration.
8. Start the VCS engine on the local node:
   `C:\> hastart`
9. Start VCS on other nodes in the cluster

Configuring the ClusterService group

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

- Cluster Management Console (Single Cluster Mode) or Web Console
- Notification
- Inter-cluster communication for global clusters

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

To configure the ClusterService group

1. Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)
2. Read the information on the Welcome panel and click **Next**.
3. In the Configuration Options panel, click **Cluster Operations** and click **Next**.
4. In the Domain Selection panel, select or type the name of the domain in which the cluster resides and click **Next**.

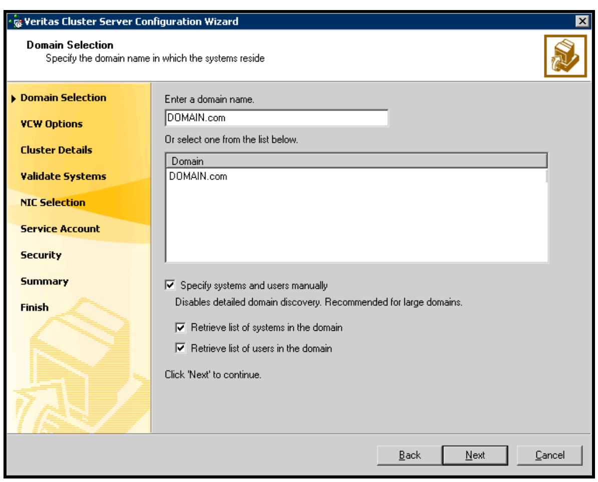

To discover information about all the systems and users in the domain

- Uncheck the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 7 on page 394.

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

- Check the **Specify systems and users manually** check box.
  Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
  If you checked the **Retrieve system list from domain** check box, proceed to step 6 on page 394. Otherwise proceed to the next step.

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

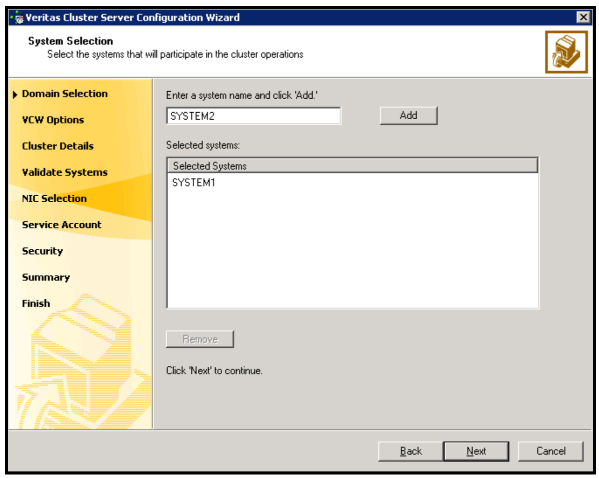

Proceed to step 7 on page 394.

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

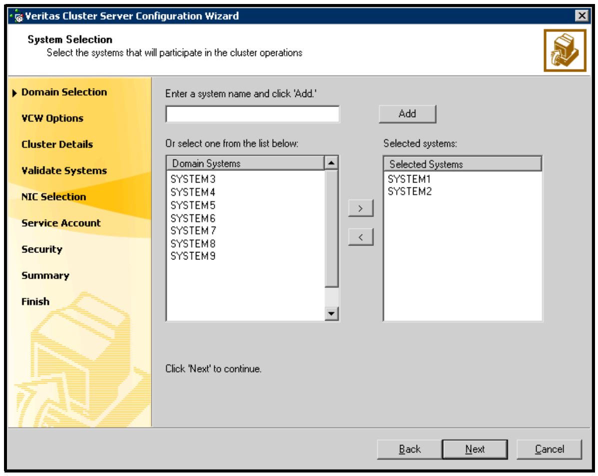

Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard will discover all the nodes for that cluster.

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

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

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

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

In the Cluster User Information dialog box, enter the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you configure a ClusterService group in a non-secure cluster.

10 In the Cluster Service Components panel, select the components to be configured in the ClusterService service group.

- Check the **Web Console** check box to configure Cluster Management Console (Single Cluster Mode) also referred to as Web Console. See “Configuring Web Console” on page 395 for instructions.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients. See “Configuring notification” on page 396 for instructions.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication. See “Configuring the wide-area connector process for global clusters” on page 400 for instructions.
- Click **Next**.

Configuring Web Console

Steps to configure the Web Console.

To configure Web Console

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

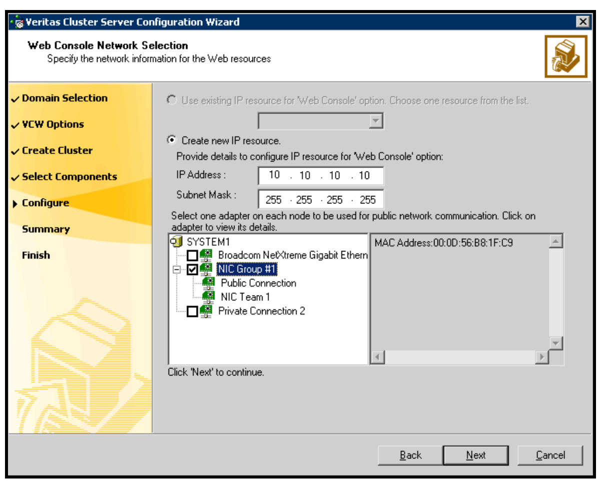

- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.

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

Configuring notification

Steps to configure notifier resource.

To configure notification

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

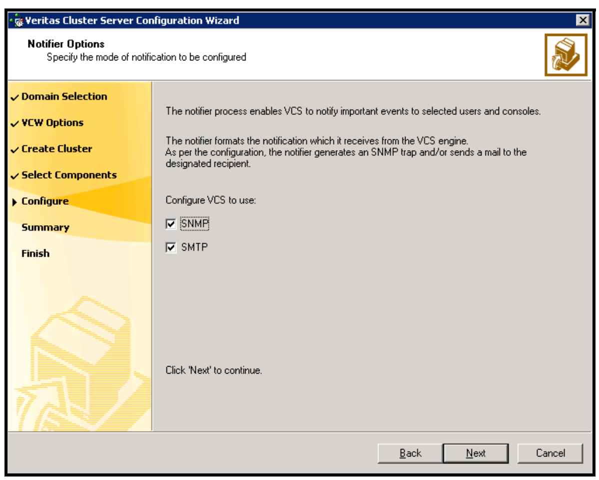

You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

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

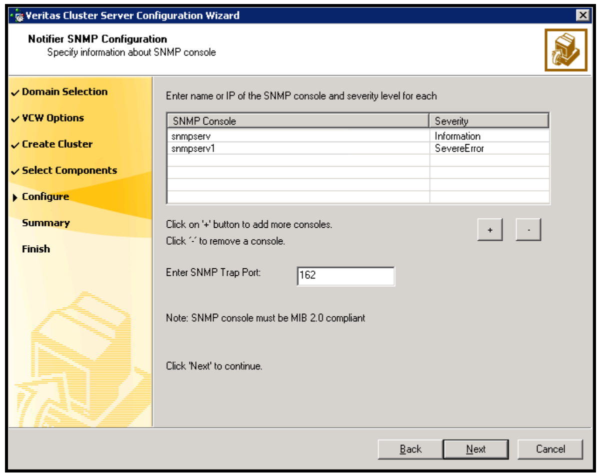

- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click **+** to add a field; click **-** to remove a field.
- 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**.

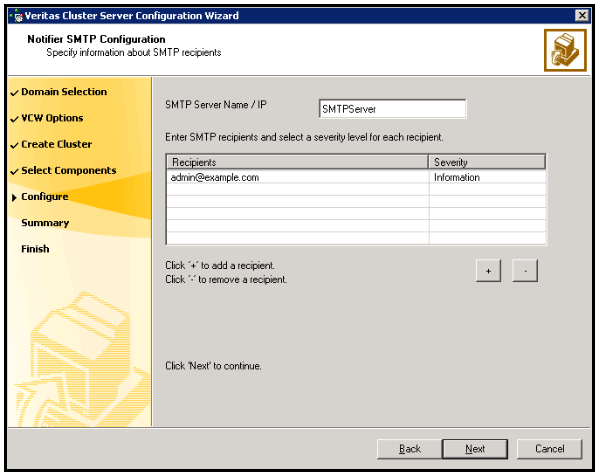

- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity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

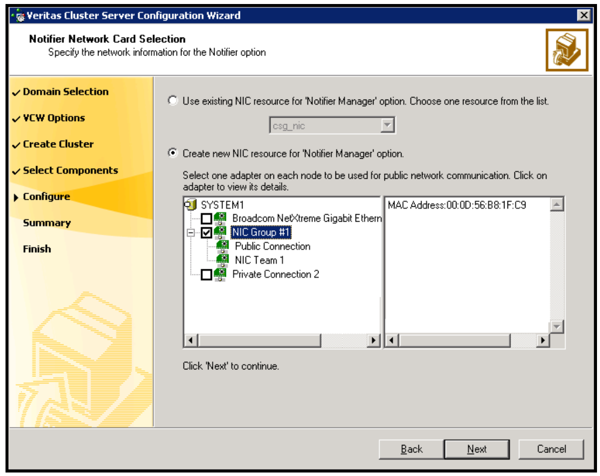

- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.

5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.

6 If you are done with the configuration, click **Finish** to exit the wizard.

Configuring the wide-area connector process for global clusters

Steps to configure wide-area connector resource for global clusters.

To configure the wide-area connector process for global clusters

1 On the GCO Network Selection panel, specify the network information and click **Next**.

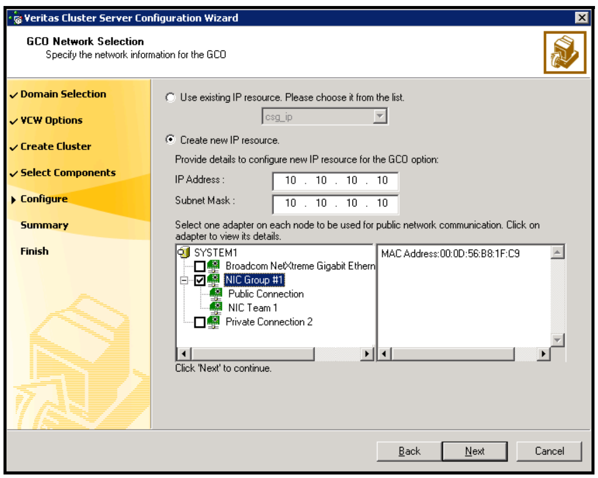

- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
- If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.

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

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

See "Setting up a global cluster" on page 553 for instructions on setting up a global cluster environment.

Deleting a cluster configuration

To delete a cluster configuration

1 Start the VCS Configuration wizard. (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Cluster Configuration Wizard**)

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

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

4 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and click **Next**.

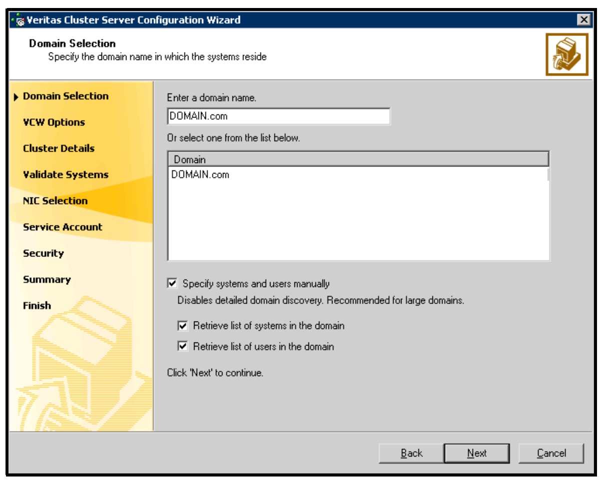

To discover information about all the systems and users in the domain

- Uncheck the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 7 on page 403.

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

- Check the **Specify systems and users manually** check box.
  Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
  If you checked the **Retrieve system list from domain** check box, proceed to step 6 on page 402. Otherwise proceed to the next step.

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

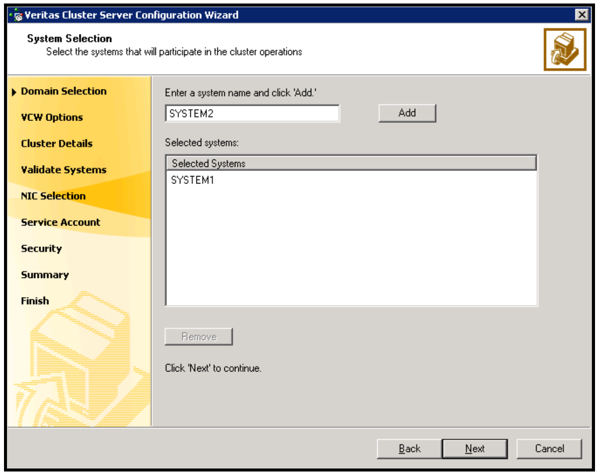

Proceed to step 7 on page 403.

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

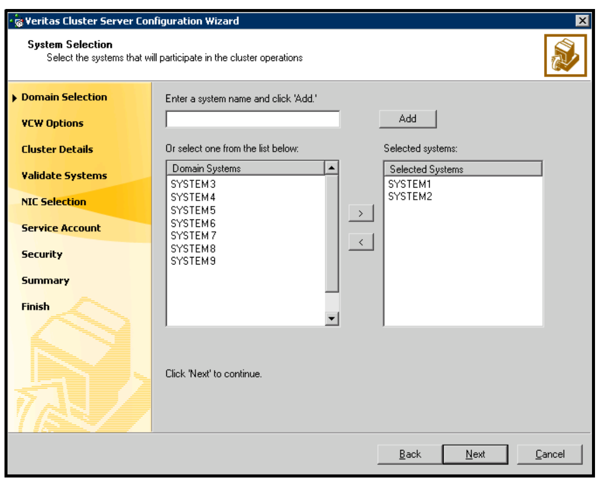

Enter the system name and click **Add** to add the system to the Selected Systems list. Alternatively, you can select the systems from the Domain Systems list and click the right-arrow icon.

If you specify only one node of an existing cluster, VCW discovers all nodes for that cluster.

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

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

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

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

10 In the Cluster User Information panel, enter the user name and password for a user with administrative privileges to the cluster and click **OK**.

   The Cluster User Information dialog box appears only when you delete a non-secure cluster.

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

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

VCW removes the cluster configuration; VCW does not unconfigure Symantec Product Authentication Service or uninstall SFW HA from the systems.

Chapter 11

Predicting VCS behavior using VCS Simulator

- About VCS Simulator
- Installing VCS Simulator
- Administering VCS Simulator from the Java Console
- Administering VCS Simulator from the command line

About VCS Simulator

VCS Simulator enables you to simulate and test cluster configurations. Use VCS Simulator to view and modify service group and resource configurations and test failover behavior. VCS Simulator can be run on a standalone system and does not require any additional hardware.

VCS Simulator runs an identical version of the VCS High Availability Daemon (HAD) as in a cluster, ensuring that failover decisions are identical to those in an actual cluster.

You can test configurations from different operating systems using VCS Simulator. For example, you can run VCS Simulator on a Windows system and test VCS configurations for Windows, Linux, and Solaris clusters. VCS Simulator also enables creating and testing global clusters.

You can administer VCS Simulator from the Java Console or from the command line.

Installing VCS Simulator

To install VCS Simulator on Windows systems

1. Insert the VCS installation disc into a drive.
2. From Windows Explorer, navigate to the path of the Simulator installer file, located at windows\VCSWindowsInstallers\Simulator\.
3. Double-click the installer file.
4. Read the information in the Welcome screen and click **Next**.
5. In the Destination Folders dialog box, click **Next** to accepted the suggested installation path or click **Change** to choose a different location.
6. In the Ready to Install the Program dialog box, click **Back** to make changes to your selections or click **Install** to proceed with the installation.
7. In the Installshield Wizard Completed dialog box, click **Finish**.

Reviewing the installation

VCS Simulator installs Cluster Manager (Java Console) and Simulator binaries on the system. The Simulator installation creates the following directories:

Directory	Contents
attrpool	Information about attributes associated with VCS objects.
bin	VCS Simulator binaries.
default_clus	Files for the default cluster configuration.
sample_clus	A sample cluster configuration, which serves as a template for each new cluster configuration.
templates	Various templates used by the Java Console.
types	The types.cf files for all supported platforms.

Additionally, VCS Simulator installs directories for various cluster configurations.

VCS Simulator creates a directory for every new simulated cluster and copies the contents of the sample_clus directory. Simulator also creates a logs directory within each cluster directory for logs associated with the cluster.

Simulator ports

VCS Simulator uses the following ports:

- Ports 15550 through 15558 to connect to the various cluster configurations.
- Ports 15560 through 15563 for the wide area connector (WAC) process.
  Set the WAC port to -1 to disable WAC simulation.

Administering VCS Simulator from the Java Console

The Simulator Console enables you to start, stop, and manage simulated clusters.

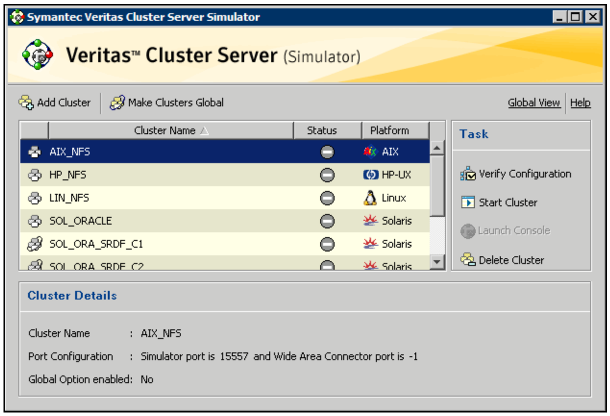

The console provides two views:

- Cluster View—Lists all simulated cluster.
- Global View—Lists global clusters.

Through the Java Console, VCS Simulator enables you to configure a simulated cluster panel, bring a system in an unknown state into an online state, simulate power loss for running systems, simulate resource faults, and save the configuration while VCS is offline. For global clusters, you can simulate the process of generating and clearing cluster faults.

You can run multiple simulated clusters on a system by using different port numbers for each cluster.

The Java Console provides the same views and features that are available for online configurations.

See “Administering the cluster from Cluster Manager (Java console)” on page 135.

Starting VCS Simulator from the Java Console

To start VCS Simulator from the Java Console (Windows)

◆ Click **Start > Programs > Symantec > VCS Simulator - Java Console**.

Creating a simulated cluster

You can start a sample cluster configuration or create a new simulated cluster.

See "Creating a simulated cluster" on page 409.

To create a simulated cluster

1. In the Simulator console, click **Add Cluster**.
2. In the Add Cluster dialog box:

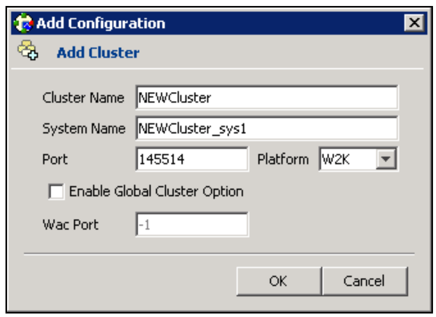

- Enter a name for the new cluster.
- Accept the suggested system name or enter a new name for a system in the cluster.
- Enter a unique port number for the simulated cluster.
- Select the platform for the cluster nodes.
- If the cluster will be part of a a global cluster configuration, select the **Enable Global Cluster Option** check box and enter a unique port number for the wide-area connector (WAC) process.
- Click **OK**.

VCS creates a simulated one-node cluster and creates a new directory for the cluster's configuration files. VCS also creates a user called *admin* with Cluster Administrator privileges. You can start the simulated cluster and administer it by launching the Java Console.

Deleting a cluster

Deleting a simulated cluster removes all files and directories associated with the cluster. Before deleting a cluster, make sure the cluster is not configured as a global cluster. You can delete global clusters from the Global View.

To delete a simulated cluster

1. From Simulator Explorer, select the cluster and click **Delete Cluster**.
2. In the confirmation dialog box, click **Yes**.

Starting a simulated cluster

Start the cluster to begin administering it.

To start a simulated cluster

1. In the Simulator console, select the cluster.
2. Click **Start Cluster**.
3. After the cluster starts, click **Launch Console** to administer the cluster.
4. Enter a valid user name and password to log on to the cluster.
   VCS Simulator does not validate passwords; you can log on to a simulated cluster by just entering a valid VCS user name. If you use the default configuration, enter admin for the user name.
   Cluster Explorer is launched upon initial logon, and the icons in the cluster panel change color to indicate an active panel.

Verifying a simulated cluster configuration

Verify the configuration is valid.

To verify the simulated cluster configuration

1. In the Simulator console, select the cluster.
2. Click **Verify Configuration**.

Simulating a global cluster configuration

Simulate a global cluster environment to test your global cluster configuration.

See "How VCS global clusters work" on page 544.

To simulate a global cluster configuration

1 Create the simulated clusters for the global configuration.
See "Creating a simulated cluster" on page 409.
Select the **Enable Global Cluster Option** check box and enter a unique port number for the wide-area connector (WAC) process.

2 In the Simulator console, click **Make Global**.

3 In the Make Global Configuration dialog box:

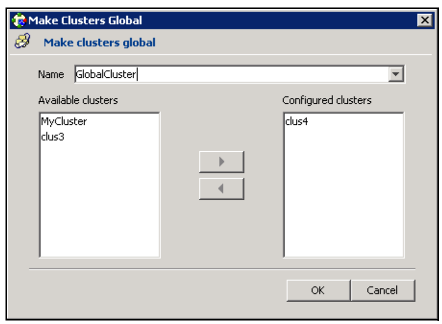

- Select an existing global cluster or enter the name for a new global cluster.
- From the **Available Clusters** list, select the clusters to add to the global cluster and click the right arrow. The clusters move to the **Configured Clusters** list.
- Click **OK**.

Bringing a system up

Bring a system up to simulate a running system

To bring a system up

1 From Cluster Explorer, click the **Systems** tab of the configuration tree.

2 Right-click the system in an unknown state, and click **Up**.

Powering off a system

1. From Cluster Explorer, click the **Systems** tab of the configuration tree.
2. Right-click the online system, and click **Power Off**.

Saving the offline configuration

1. From Cluster Explorer, click **Save Configuration As** from the **File** menu.
2. Enter the path location.
3. Click **OK**.

Simulating a resource fault

1. From Cluster Explorer, click the **Service Groups** tab of the configuration tree.
2. Right-click an online resource, click **Fault Resource**, and click the system name.

Simulating cluster faults in global clusters

Use VCS Simulator to imitate the process of generating and clearing cluster faults.

See "Monitoring alerts" on page 232.

To generate a cluster fault

1. From Cluster Explorer, click the cluster in the configuration tree.
2. Right-click the cluster, click **Fault Cluster**, and click the cluster name.

To clear a cluster fault

1. From Cluster Explorer, click the cluster in the configuration tree.
2. Right-click the cluster, click **Clear Cluster Fault**, and click the cluster name.

Administering VCS Simulator from the command line

Start VCS Simulator before creating or administering simulated clusters.

Note: VCS Simulator treats clusters created from the command line and the Java Console separately. So, clusters created from the command line are not visible in the graphical interface. Also, If you delete a cluster from the command line, you may see the cluster in the Java Console.

Starting VCS Simulator from the command line

To start VCS Simulator from the command line (Windows)

VCS Simulator installs platform-specific types.cf files at the path *%VCS_SIMULATOR_HOME%*\types\. The variable *%VCS_SIMULATOR_HOME%* represents the Simulator installation directory, typically C:\Program Files\Veritas\VCS Simulator\ .

1. To simulate a cluster running a particular operating system, copy the types.cf. file for the operating system from the types directory to *%VCS_SIMULATOR_HOME%*\default_clus\conf\config\.
   For example, if the cluster to be simulated runs on the AIX platform, copy the file types.cf.aix.
2. Add custom type definitions to the file, if required, and rename the file to types.cf.
3. If you have a main.cf file to run in the simulated cluster, copy it to *%VCS_SIMULATOR_HOME%*\default_clus\conf\config\.
4. Start VCS Simulator:
   `%VCS_SIMULATOR_HOME%\bin>` **`hasim -start`** ***`system_name`***
   The variable *system_name* represents a system name, as defined in the configuration file main.cf.
   This command starts Simulator on port 14153.
5. Add systems to the configuration, if desired:
   `%VCS_SIMULATOR_HOME%\bin>` **`hasim -sys -add`** ***`system_name`***
   `%VCS_SIMULATOR_HOME%\bin>` **`hasim -up`** ***`system_name`***
6. Verify the state of each node in the cluster:
   `%VCS_SIMULATOR_HOME%\bin>` **`hasim -sys -state`**
   See "To simulate global clusters from the command line" on page 414.

To simulate global clusters from the command line

1 Install VCS Simulator in a directory (*sim_dir*) on your system.
See"Installing VCS Simulator" on page 406.

2 Set up the clusters on your system. Run the following command to add a cluster:

```
sim_dir/hasim -setupclus clustername -simport
   port_no -wacport port_no
```

Do not use *default_clus* as the cluster name when simulating a global cluster.

VCS Simulator copies the sample configurations to the path sim_dir/ clustername and creates a system named `clustername_sys1`.

For example, to add cluster clus_a using ports 15555 and 15575, run the following command:

```
sim_dir/hasim -setupclus clus_a -simport 15555 -wacport 15575
```

Similarly, add the second cluster:

```
sim_dir/hasim -setupclus clus_b -simport 15556 -wacport 15576
```

To create multiple clusters without simulating a global cluster environment, specify -1 for the wacport.

3 Start the simulated clusters:

```
sim_dir/hasim -start clustername_sys1 -clus clustername
```

4 Set the following environment variables to access VCS Simulator from the command line:

- VCS_SIM_PORT=*port_number*
- VCS_SIM_WAC_PORT=*wacport*

Note that you must set these variables for each simulated cluster, otherwise Simulator always connects default_clus, the default cluster.

You can use the Java Console to link the clusters and to configure global service groups.

See"Administering the cluster from Cluster Manager (Java console)" on page 135.

You can also edit the configuration file main.cf manually to create the global cluster configuration.

Administering simulated clusters from the command line

The functionality of VCS Simulator commands mimic that of standard `ha` commands.

Command	Description
`hasim -start` *system_name*	Starts VCS Simulator. The variable *system_name* represents the system that will transition from the LOCAL_BUILD state to RUNNING.
`hasim -setupclus` *clustername* `-simport port_no [-wacport port_no] [-sys systemname]`	Creates a simulated cluster and associates the specified ports with the cluster.
`hasim -deleteclus <clus>`	Deletes specified cluster. Deleting the cluster removes all files and directories associated with the cluster. Before deleting a cluster, make sure the cluster is not configured as a global cluster.
`hasim -start` *clustername_sys1* `[-clus` *clustername*`] [-disablel10n]`	Starts VCS Simulator on the cluster specified by *clustername*. If you start VCS Simulator with the `-disablel10n` option, the simulated cluster does not accept localized values for attributes. Use this option when simulating a UNIX configuration on a Windows system to prevent potential corruption when importing the simulated configuration to a UNIX cluster.
`hasim -stop`	Stops the simulation process.
`hasim -poweroff` *system_name*	Gracefully shuts down the system.
`hasim -up` *system_name*	Brings the system up.
`hasim -fault` *system_name resource_name*	Faults the specified resource on the specified system.
`hasim -online system_name` *resource_name*	Brings specified resource online. This command is useful if you have simulated a fault of a persistent resource and want to simulate the fix.
`hasim -faultcluster` *clustername*	Simulates a cluster fault.
`hasim -clearcluster` *clustername*	Clears a simulated cluster fault.

Command	Description
`hasim -getsimconfig` *`cluster_name`*	Retrieves information about VCS Simulator ports.
`hasim -hb [..]`	Equivalent to standard `hahb` command.
`hasim -disablel10n`	Disables localized inputs for attribute values. Use this option when simulating UNIX configurations on Windows systems.
`hasim -clus [...]`	Equivalent to standard `haclus` command.
`hasim -sys [...]`	Equivalent to standard `hasys` command.
`hasim -grp [...]`	Equivalent to standard `hagrp` command.
`hasim -res [...]`	Equivalent to standard `hares` command.
`hasim -type [...]`	Equivalent to standard `hatype` command.
`hasim -conf [...]`	Equivalent to standard `haconf` command.
`hasim -attr [...]`	Equivalent to standard `haattr` command.

Section III

VCS communication and operations

- Chapter 12, "About communications, membership, and data protection in the cluster" on page 419
- Chapter 13, "Controlling VCS behavior" on page 447
- Chapter 14, "The role of service group dependencies" on page 491

About communications, membership, and data protection in the cluster

- About cluster communications
- About cluster membership
- About membership arbitration
- About data protection
- Examples of VCS operation with I/O fencing
- About cluster membership and data protection without I/O fencing
- Examples of VCS operation without I/O fencing
- Summary of best practices for cluster communications

About cluster communications

VCS uses local communications on a system and system-to-system communications.

About intra-system communications

Within a system, the VCS engine (HAD) uses a VCS-specific communication protocol known as Inter Process Messaging (IPM) to communicate with the GUI, the command line, and the agents. The following illustration shows basic communication on a single VCS system. Note that agents only communicate with HAD and never communicate with each other.

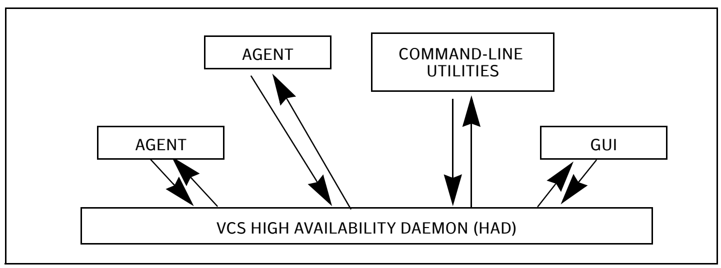

The following illustration depicts communication from a single agent to HAD.

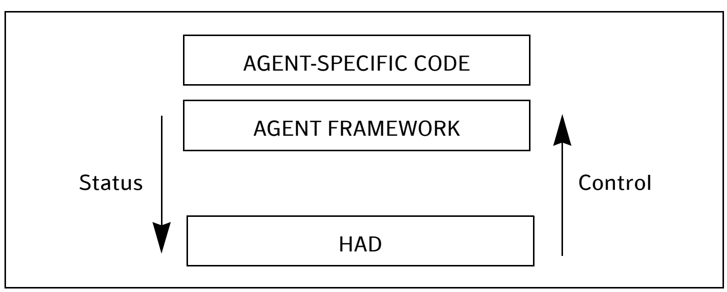

The agent uses the agent framework, which is compiled into the agent itself. For each resource type configured in a cluster, an agent runs on each cluster system. The agent handles all resources of that type. The engine passes commands to the agent and the agent returns the status of command execution. For example, an agent is commanded to bring a resource online. The agent responds back with the success (or failure) of the operation. Once the resource is online, the agent communicates with the engine only if this status changes.

About inter-system cluster communications

VCS uses the cluster interconnect for network communications between cluster systems. Each system runs as an independent unit and shares information at the cluster level. On each system the VCS High Availability Daemon (HAD), which is the decision logic for the cluster, maintains a view of the cluster configuration. This daemon operates as a replicated state machine, which means all systems in the cluster have a synchronized state of the cluster configuration. This is accomplished by the following:

- All systems run an identical copy of HAD.
- HAD on each system maintains the state of its own resources, and sends all cluster information about the local system to all other machines in the cluster.
- HAD on each system receives information from the other cluster systems to update its own view of the cluster.
- Each system follows the same code path for actions on the cluster.

The replicated state machine communicates over a purpose-built communications package consisting of two components, *Group Membership Services/Atomic Broadcast (GAB)* and *Low Latency Transport (LLT)*.

Figure 12-13 illustrates the overall communications paths between two systems of the replicated state machine model.

Figure 12-13 Cluster communications with replicated state machine

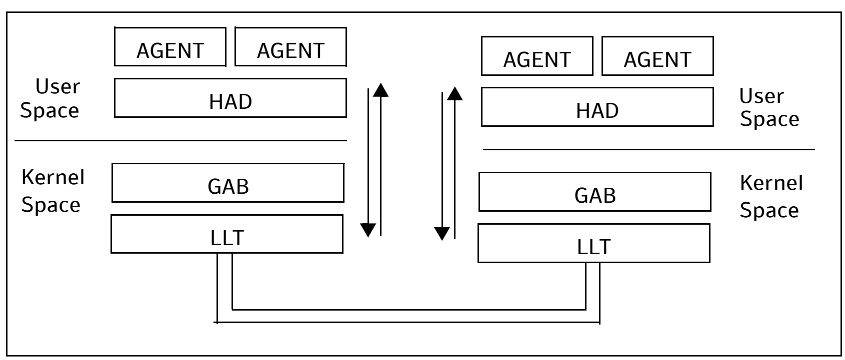

Group Membership Services/Atomic Broadcast (GAB)

The Group Membership Services/Atomic Broadcast protocol (GAB) is responsible for cluster membership and reliable cluster communications. GAB has two major functions.

- Cluster membership
  GAB maintains cluster membership by receiving input on the status of the heartbeat from each system via LLT. When a system no longer receives heartbeats from a cluster peer, LLT passes the heartbeat loss to GAB. GAB marks the peer as DOWN and excludes it from the cluster. In most configurations, membership arbitration is used to prevent network partitions.
- Cluster communications
  GAB's second function is reliable cluster communications. GAB provides guaranteed delivery of messages to all cluster systems. The Atomic Broadcast functionality is used by HAD to ensure that all systems within the cluster receive all configuration change messages, or are rolled back to the previous state, much like a database atomic commit. While the communications function in GAB is known as Atomic Broadcast, no actual network broadcast traffic is generated. An Atomic Broadcast message is a series of point to point unicast messages from the sending system to each receiving system, with a corresponding acknowledgement from each receiving system.

Low Latency Transport (LLT)

The Low Latency Transport protocol is used for all cluster communications as a high-performance, low-latency replacement for the IP stack. LLT has two major functions.

- Traffic distribution
  LLT provides the communications backbone for GAB. LLT distributes (load balances) inter-system communication across all configured network links. This distribution ensures all cluster communications are evenly distributed across all network links for performance and fault resilience. If a link fails, traffic is redirected to the remaining links. A maximum of eight network links are supported.
- Heartbeat
  LLT is responsible for sending and receiving heartbeat traffic over each configured network link. LLT heartbeat is an Ethernet broadcast packet. This broadcast heartbeat method allows a single packet to notify all other cluster members the sender is functional, as well as provide necessary address information for the receiver to send unicast traffic back to the

sender. The heartbeat is the only broadcast traffic generated by VCS. Each system sends 2 heartbeat packets per second per interface. All other cluster communications, including all status and configuration traffic is point to point unicast. This heartbeat is used by the Group Membership Services to determine cluster membership.

The heartbeat signal is defined as follows:

- LLT on each system in the cluster sends heartbeat packets out on all configured LLT interfaces every half second.
- LLT on each system tracks the heartbeat status from each peer on each configured LLT interface.
- LLT on each system forwards the heartbeat status of each system in the cluster to the local Group Membership Services function of GAB.
- GAB receives the status of heartbeat from all cluster systems from LLT and makes membership determination based on this information.

Figure 12-14 Heartbeat in the cluster

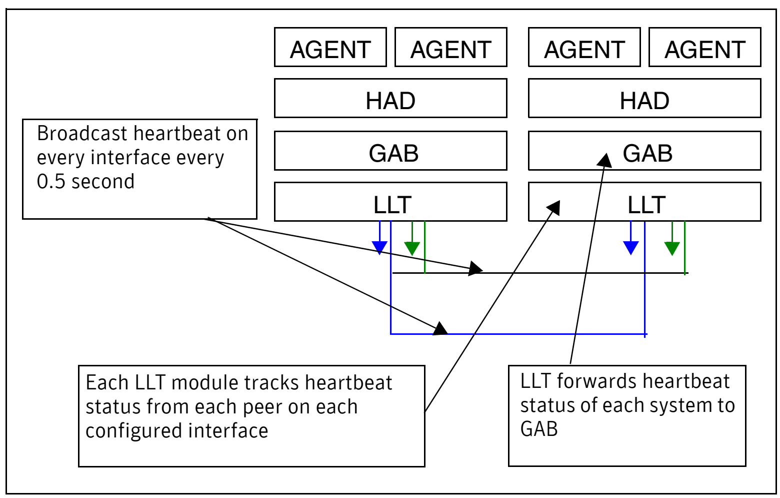

LLT can be configured to designate specific cluster interconnect links as either high priority or low priority. High priority links are used for cluster communications to GAB as well as heartbeat signals. Low priority links, during normal operation, are used for heartbeat and link state maintenance only, and the frequency of heartbeats is reduced to 50% of normal to reduce network overhead.

If there is a failure of all configured high priority links, LLT will switch all cluster communications traffic to the first available low priority link. Communication traffic will revert back to the high priority links as soon as they become available.

While not required, best practice recommends to configure at least one low priority link, and to configure two high priority links on dedicated cluster interconnects to provide redundancy in the communications path. Low priority links are typically configured on the public or administrative network.

About cluster membership

The current members of the cluster are the systems that are actively participating in the cluster. It is critical for HAD to accurately determine current cluster membership in order to take corrective action on system failure and maintain overall cluster topology.

A change in cluster membership is one of the starting points of the logic to determine if HAD needs to perform any fault handling in the cluster.

There are two aspects to cluster membership, initial joining of the cluster and how membership is determined once the cluster is up and running.

Initial joining of systems to cluster membership

When the cluster initially boots, LLT determines which systems are sending heartbeat signals, and passes that information to GAB. GAB uses this information in the process of seeding the cluster membership.

Seeding a new cluster

Seeding insures a new cluster will start with an accurate membership count of the number of systems in the cluster. This prevents the possibility of one cluster splitting into multiple subclusters upon initial startup. A new cluster can be automatically seeded as follows:

- When the cluster initially boots, all systems in the cluster are unseeded.
- GAB checks the number of systems that have been declared to be members of the cluster in the /etc/gabtab file.

  The number of systems declared in the cluster is denoted as follows:

  ```
  /sbin/gabconfig -c -n#
  ```

  where the variable # is replaced with the number of systems in the cluster.

  Note: This number should represent 100% of the systems in the cluster

- When GAB on each system detects that the correct number of systems are running, based on the number declared in /etc/gabtab and input from LLT, it will seed.
- HAD will start on each seeded system. HAD will only run on a system that has seeded.

Manual seeding of a cluster

Seeding the cluster manually is appropriate when the number of cluster systems declared in /etc/gabtab is more than the number of systems that will join the cluster. This could occur if a system is down for maintenance when the cluster comes up.

Caution: It is not recommended to seed the cluster manually unless the administrator is aware of the risks and implications of the command.

Before manually seeding the cluster, check that systems that will join the cluster are able to send and receive heartbeats to each other. Confirm there is no possibility of a network partition condition in the cluster.

To manually seed the cluster, type the following command:

```
/sbin/gabconfig -c -x
```

Note there is no declaration of the number of systems in the cluster with a manual seed. This command will seed all systems in communication with the system where the command is run.

See “Seeding and I/O Fencing” on page 538.

Ongoing cluster membership

Once the cluster is up and running, a system remains an active member of the cluster as long as peer systems receive a heartbeat signal from that system over the cluster interconnect. A change in cluster membership is determined as follows:

- When LLT on a system no longer receives heartbeat messages from a system on any of the configured LLT interfaces for a predefined time, LLT informs GAB of the heartbeat loss from that specific system.
  This predefined time is 16 seconds by default, but can be configured. It is set with the `set-timer peerinact` command as described in the llttab manual page.
- When LLT informs GAB of a heartbeat loss, the systems that are remaining in the cluster coordinate to agree which systems are still actively participating in the cluster and which are not. This happens during a time period known as GAB Stable Timeout (5 seconds).
  VCS has specific error handling that takes effect in the case where the systems do not agree.
- GAB marks the system as DOWN, excludes the system from the cluster membership, and delivers the membership change to the fencing module.
- The fencing module performs membership arbitration to ensure that there is not a split brain situation and only one functional cohesive cluster continues to run.

The fencing module is turned on by default.
See "About cluster membership and data protection without I/O fencing" for actions that occur if the fencing module has been deactivated.

About membership arbitration

Membership arbitration is necessary on a perceived membership change because systems may falsely appear to be down. When LLT on a system no longer receives heartbeat messages from another system on any configured LLT interface, GAB marks the system as DOWN. However, if the cluster interconnect network failed, a system can appear to be failed when it actually is not. In most environments when this happens, it is caused by an insufficient cluster interconnect network infrastructure, usually one that routes all communication links through a single point of failure.

If all the cluster interconnect links fail, it is possible for one cluster to separate into two subclusters, each of which does not know about the other subcluster. The two subclusters could each carry out recovery actions for the departed systems. This is termed split brain.

In a split brain condition, two systems could try to import the same storage and cause data corruption, have an IP address up in two places, or mistakenly run an application in two places at once.

Membership arbitration guarantees against such split brain conditions.

Components of membership arbitration

The components of membership arbitration are the fencing module and the coordinator disks.

Fencing module

Each system in the cluster runs a kernel module called vxfen, or the *fencing module*. This module is responsible for ensuring valid and current cluster membership on a membership change through the process of membership arbitration. vxfen performs the following actions:

- Registers with the coordinator disks during normal operation
- Races for control of the coordinator disks during membership changes

Coordinator disks

Coordinator disks are a number of special purpose disks that act together as a global lock device. Racing for control of these disks is used to determine cluster membership. Control is won by the system that gains control of a majority of the coordinator disks, so there must always be an odd number of disks, with three disks recommended.

Coordinator disks cannot be used for any other purpose in the cluster configuration, such as data storage or inclusion in a disk group for user data.

Any disks that support SCSI-3 Persistent Reservation can be coordinator disks. Best practice is to select the smallest possible LUNs for use as coordinator disks.

You can configure coordinator disks to use Veritas Volume Manager Dynamic Multipathing (DMP) feature. For more information on using DMP, see the *Veritas Volume Manager Administrator's* Guide.

How the fencing module starts up

The fencing module starts up as follows:

- The coordinator disks are placed in a disk group.
  This allows the fencing start up script to use Veritas Volume Manager (VxVM) commands to easily determine which disks are coordinator disks, and what paths exist to those disks. This disk group is never imported, and is not used for any other purpose.
- The fencing start up script on each system uses VxVM commands to populate the file /etc/vxfentab with the paths available to the coordinator disks.
  For example, if the user has configured 3 coordinator disks with 2 paths to each disk, the /etc/vxfentab file will contain 6 individual lines, representing the path name to each disk, such as
- When the fencing driver is started, it reads the physical disk names from the /etc/vxfentab file. Using these physical disk names, it determines the serial numbers of the coordinator disks and builds an in-memory list of the drives.
- The fencing driver verifies that any other systems in the cluster that are already up and running see the same coordinator disks.
  The fencing driver examines GAB port B for membership information. If no other systems are up and running, it is the first system up and is considered the correct coordinator disk configuration. When a new member joins, it requests a coordinator disks configuration. The system with the lowest LLT ID will respond with a list of the coordinator disk serial numbers. If there is a match, the new member joins the cluster. If there is not a match, vxfen enters an error state and the new member is not allowed to join. This process ensures all systems communicate with the same coordinator disks.
- The fencing driver determines if a possible preexisting split brain condition exists.
  This is done by verifying that any system that has keys on the coordinator disks can also be seen in the current GAB membership. If this verification fails, the fencing driver prints a warning to the console and system log and does not start.
- If all verifications pass, the fencing driver on each system registers keys with each coordinator disk.

Figure 12-15 Topology of coordinator disks in the cluster

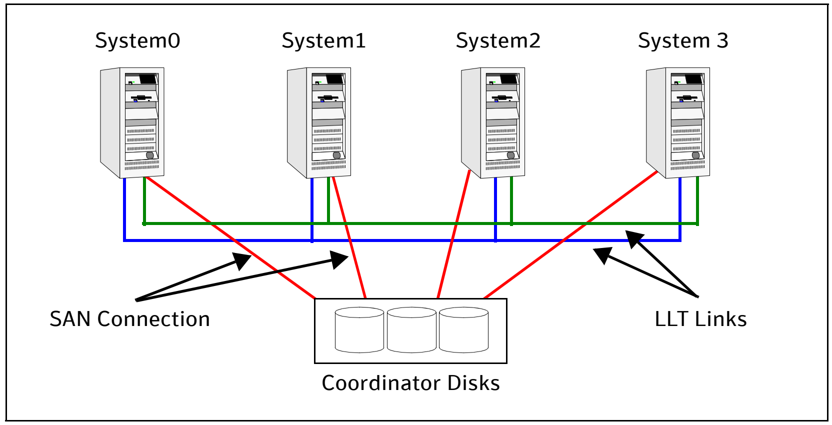

How membership arbitration works

Upon startup of the cluster, all systems register a unique key on the coordinator disks. (The key is based on the LLT system ID, for example LLT ID 0 = A.) When there is a perceived change in membership, membership arbitration works as follows:

- GAB marks the system as DOWN, excludes the system from the cluster membership, and delivers the membership change—the list of departed systems—to the fencing module.
- The system with the lowest LLT system ID in the cluster races for control of the coordinator disks
  - In the most common case, where departed systems are truly down or faulted, this race has only one contestant.
  - In a split brain scenario, where two or more subclusters have formed, the race for the coordinator disks is performed by the system with the lowest LLT system ID of that subcluster. This system races on behalf of all the other systems in its subcluster.
- The race consists of executing a preempt and abort command for each key of each system that appears to no longer be in the GAB membership.

  The preempt and abort command allows only a registered system with a valid key to eject the key of another system. This ensures that even when multiple systems attempt to eject other, each race will have only one winner. The first system to issue a preempt and abort command will win and eject the key of the other system. When the second system issues a preempt and abort command, it can not perform the key eject because it is no longer a registered system with a valid key.

- If the preempt and abort command returns success, that system has won the race for that coordinator disk.
  Each system will repeat this race to all the coordinator disks. The race is won by, and control is attained by, the system that ejects the other system's registration keys from a majority of the coordinator disks.
- On the system that wins the race, the vxfen module informs all the systems that it was racing on behalf of that it won the race, and that subcluster is still valid. This information is passed back to GAB.
- On the system(s) that do not win the race, the vxfen module will trigger a system panic. The other systems in this subcluster will note the panic, determine they lost control of the coordinator disks, and also panic and restart.
- Upon restart, the systems will attempt to seed into the cluster.
  - If the systems that restart can exchange heartbeat with the number of cluster systems declared in /etc/gabtab, they will automatically seed and continue to join the cluster. Their keys will be replaced on the coordinator disks. This case will only happen if the original reason for the membership change has cleared during the restart.
  - If the systems that restart can not exchange heartbeat with the number of cluster systems declared in /etc/gabtab, they will not automatically seed, and HAD will not start. This is a possible split brain condition, and requires administrative intervention.

Note: Forcing a manual seed at this point will allow the cluster to seed. However, when the fencing module checks the GAB membership against the systems that have keys on the coordinator disks, a mismatch will occur. vxfen will detect a possible split brain condition, print a warning, and will not start. In turn, HAD will not start. Administrative intervention is required.

About data protection

Membership arbitration by itself is inadequate for complete data protection because it assumes that all systems will either participate in the arbitration or are already down.

Rare situations can arise which must also be protected against. Some examples are:

- A system hang causes the kernel to stop processing for a period of time.
- The system resources were so busy that the heartbeat signal was not sent.
- A break and resume function is supported by the hardware and executed. Dropping the system to a system controller level with a break command can result in the heartbeat signal timeout.

In these types of situations, the systems are not actually down, and may return to the cluster after cluster membership has been recalculated. This could result in data corruption as a system could potentially write to disk before it determines it should no longer be in the cluster.

Combining membership arbitration with data protection of the shared storage eliminates all of the above possibilities for data corruption.

Data protection fences off (removes access to) the shared data storage from any system that is not a current and verified member of the cluster. Access is blocked by the use of SCSI-3 persistent reservations.

SCSI-3 Persistent Reservation

SCSI-3 Persistent Reservation (SCSI-3 PR) supports device access from multiple systems, or from multiple paths from a single system. At the same time it blocks access to the device from other systems, or other paths.

VCS logic determines when to online a service group on a particular system. If the service group contains a disk group, the disk group is imported as part of the service group being brought online. When using SCSI-3 PR, importing the disk group puts registration and reservation on the data disks. Only the system that has imported the storage with SCSI-3 reservation can write to the shared storage. This prevents a system that did not participate in membership arbitration from corrupting the shared storage.

SCSI-3 PR ensures persistent reservations across SCSI bus resets.

Note: Use of SCSI 3 PR protects against all elements in the IT environment that might be trying to write illegally to storage, not only VCS related elements.

Membership arbitration combined with data protection is termed I/O Fencing.

Examples of VCS operation with I/O fencing

This topic describes the general logic employed by the I/O fencing module along with some specific example scenarios.

About the I/O fencing algorithm

To ensure the most appropriate behavior is followed in both common and rare corner case events, the fencing algorithm works as follows:

- The fencing module is designed to never have systems in more than one subcluster remain current and valid members of the cluster. In all cases, either one subcluster will survive, or in very rare cases, no systems will.
- The system with the lowest LLT ID in any subcluster of the original cluster races for control of the coordinator disks on behalf of the other systems in that subcluster.
- If a system wins the race for the first coordinator disk, that system is given priority to win the race for the other coordinator disks.
  Any system that loses a race will delay a short period of time before racing for the next disk. Under normal circumstances, the winner of the race to the first coordinator disk will win all disks.
  This ensures a clear winner when multiple systems race for the coordinator disk, preventing the case where three or more systems each win the race for one coordinator disk.
- If the cluster splits such that one of the subclusters has at least 51% of the members of the previous stable membership, that subcluster is given priority to win the race.
  The system in the smaller subcluster(s) delay a short period before beginning the race.
  This ensures that as many systems as possible will remain running in the cluster.
- If the vxfen module discovers on startup that the system that has control of the coordinator disks is not in the current GAB membership, an error message indicating a possible split brain condition is printed to the console.
  The administrator must clear this condition manually with the `vxfenclearpre` utility.

Two system cluster where one system fails

In this example, System1 fails, and System0 carries out the I/O fencing operation as follows:

- The GAB module on System0 determines System1 has failed due to loss of heartbeat signal reported from LLT.
- GAB passes the membership change to the fencing module on each system in the cluster.
  The only system that is still running is System0
- System0 gains control of the coordinator disks by ejecting the key registered by System1 from each coordinator disk.
  The ejection takes place one by one, in the order of the coordinator disk's serial number.
- When the fencing module on System0 successfully controls the coordinator disks, HAD carries out any associated policy connected with the membership change.
- System1 is blocked access to the shared storage, if this shared storage was configured in a service group that was now taken over by System0 and imported.

Figure 12-16 I/O Fencing example with system failure

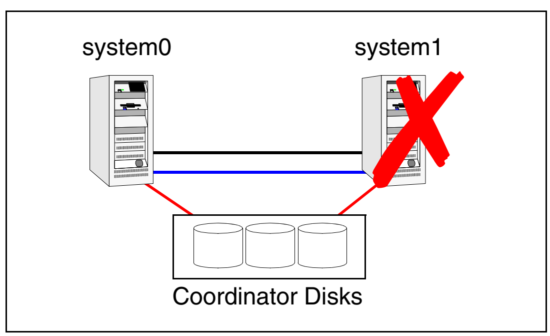

Four system cluster where cluster interconnect fails

In this example, the cluster interconnect fails in such a way as to split the cluster from one four-system cluster to two-system clusters. The cluster performs membership arbitration to ensure that only one subcluster remains.

Due to loss of heartbeat, System0 and System1 both believe System2 and System3 are down. System2 and System3 both believe System0 and System1 are down. The progression of I/O fencing operations are as follows:

Figure 12-17 Four system cluster where cluster interconnect fails

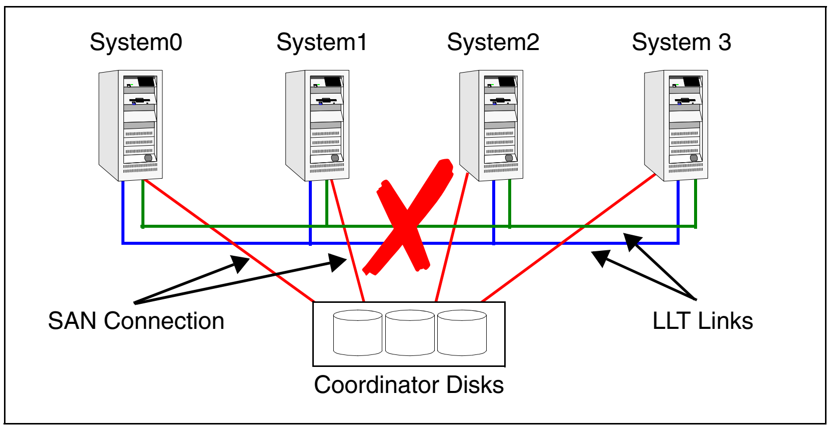

- LLT on each of the four systems no longer receives heartbeat messages from the systems on the other side of the interconnect failure on any of the configured LLT interfaces for the peer inactive timeout configured time.
- LLT on each machine passes to GAB that it has noticed a membership change. Specifically:
  - LLT on System0 passes to GAB that it no longer sees System2 and System3
  - LLT on System1 passes to GAB that it no longer sees System2 and System3
  - LLT on System2 passes to GAB that it no longer sees System0 and System1
  - LLT on System3 passes to GAB that it no longer sees System0 and System1
- After LLT informs GAB of a heartbeat loss, the systems that are remaining do a “GAB Stable Timeout (5 seconds). In this example:

- System0 and System1 agree that both of them do not see System2 and System3
- System2 and System3 agree that both of them do not see System0 and System1

■ GAB marks the system as DOWN, and excludes the system from the cluster membership. In this example:
  - GAB on System0 and System1 mark System2 and System3 as DOWN and excludes them from cluster membership.
  - GAB on System2 and System3 mark System0 and System1 as DOWN and excludes them from cluster membership.

■ GAB on each of the four systems passes the membership change to the vxfen driver for membership arbitration. Each subcluster races for control of the coordinator disks. In this example:
  - System0 has the lower LLT ID, and races on behalf of itself and System1.
  - System2 has the lower LLT ID, and races on behalf of itself and System3.

■ GAB on each of the four systems also passes the membership change to HAD. HAD waits for the result of the membership arbitration from the fencing module before taking any further action.

■ Assume System0 wins the race for the coordinator disks, and ejects the registration keys of System2 and System3 off the disks. The result is as follows:
  - System0 wins the race for the coordinator disk. The fencing module on System0 communicates race success to all other fencing modules in the current cluster, in this case System0 and System1. The fencing module on each system in turn communicates success to HAD. System0 and System1 remain valid and current members of the cluster.
  - System2 loses the race for control of the coordinator disks. The fencing module on System2 calls a kernel panic and the system restarts.
  - System3 sees another membership change from the kernel panic of System2. Because that was the system that was racing for control of the coordinator disks in this subcluster, System3 also panics.

■ HAD carries out any associated policy or recovery actions based on the membership change.

■ System2 and System3 are blocked access to the shared storage (if the shared storage was part of a service group that is now taken over by System0 or System 1).

- To rejoin System2 and System3 to the cluster, the administrator must do the following:
  - Shut down System2 and System3
  - Fix the cluster interconnect links
  - Restart System2 and System3

About cluster membership and data protection without I/O fencing

Proper seeding of the cluster and the use of low priority heartbeat cluster interconnect links are best practices with or without the use of I/O fencing. Best practice also recommends multiple cluster interconnect links between systems in the cluster. This allows GAB to differentiate between:

- A loss of all heartbeat links simultaneously, which is interpreted as a system failure. In this case, depending on failover configuration, HAD may attempt to restart the services that were running on that system on another system.
- A loss of all heartbeat links over time, which is interpreted as an interconnect failure. In this case, the assumption is made that there is a high probability that the system is not down, and HAD does not attempt to restart the services on another system.

In order for this differentiation to have meaning, it is important to ensure the cluster interconnect links do not have a single point of failure, such as a network hub or ethernet card.

Jeopardy

In all cases, when LLT on a system no longer receives heartbeat messages from another system on any of the configured LLT interfaces, GAB reports a change in membership.

When a system has only one interconnect link remaining to the cluster, GAB can no longer reliably discriminate between loss of a system and loss of the network. The reliability of the system's membership is considered at risk. A special membership category takes effect in this situation, called a jeopardy membership. This provides the best possible split-brain protection without membership arbitration and SCSI-3 capable devices.

When a system is placed in jeopardy membership status, two actions occur

- Service groups running on the system are placed in autodisabled state. A service group in autodisabled state may failover on a resource or group fault, but can not fail over on a system fault until the autodisabled flag is manually cleared by the administrator.
- VCS operates the system as a single system cluster. Other systems in the cluster are partitioned off in a separate cluster membership.

Daemon Down Node Alive (DDNA)

Daemon Down Node Alive (DDNA) is a condition in which the VCS high availability daemon (HAD) on a node fails, but the node is running. When HAD fails, the hashadow process tries to bring HAD up again. If the hashadow process succeeds in bringing HAD up, the system leaves the DDNA membership and joins the regular membership.

In a DDNA condition, VCS does not have information about the state of service groups on the node. So, VCS places all service groups that were online on the affected node in the autodisabled state. The service groups that were online on the node cannot fail over.

Manual intervention is required to enable failover of autodisabled service groups. The administrator must release the resources running on the affected node, clear resource faults, and bring the service groups online on another node.

Examples of VCS operation without I/O fencing

The following scenarios describe events, and how VCS responds, in a cluster without I/O fencing.

Four system cluster without a low priority link

Consider a four-system cluster that has two private cluster interconnect heartbeat links. The cluster does not have any low priority link.

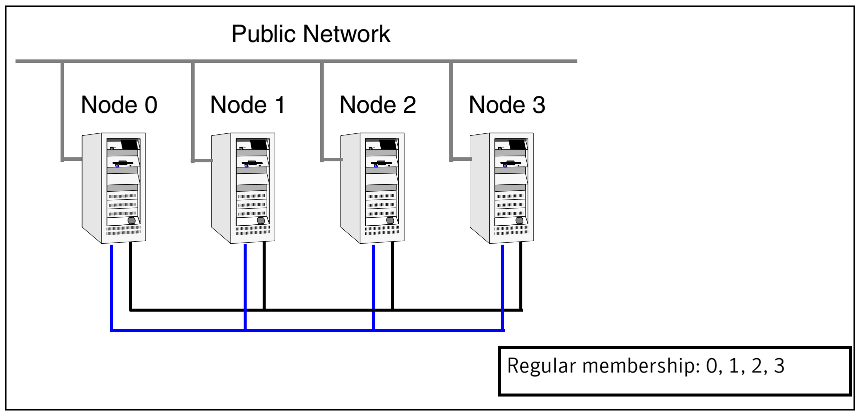

Cluster interconnect link failure

In this example, a link to System2 fails, leaving System2 with only one cluster interconnect link remaining.

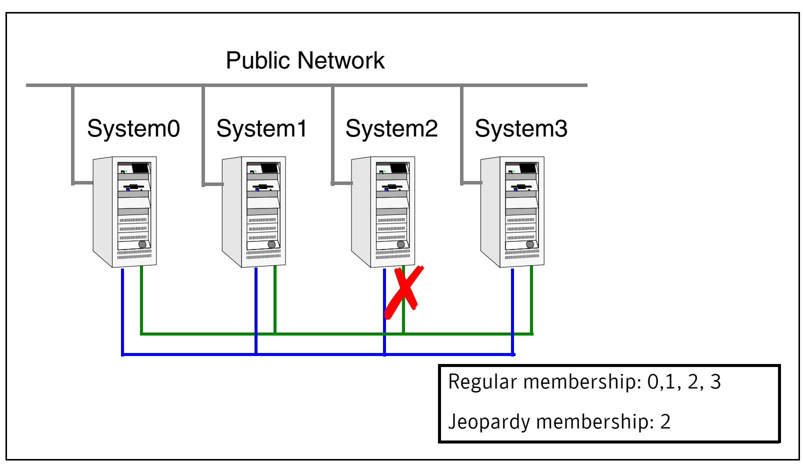

The cluster is reformed. Systems 0, 1, 2, and 3 are in the regular membership and System2 in a jeopardy membership. Service groups on System2 are autodisabled. All normal cluster operations continue, including normal failover of service groups due to resource fault.

Cluster interconnect link failure followed by system failure

In this example, the link to System2 fails, and System2 is put in the jeopardy membership. Subsequently, System2 fails due to a power fault.

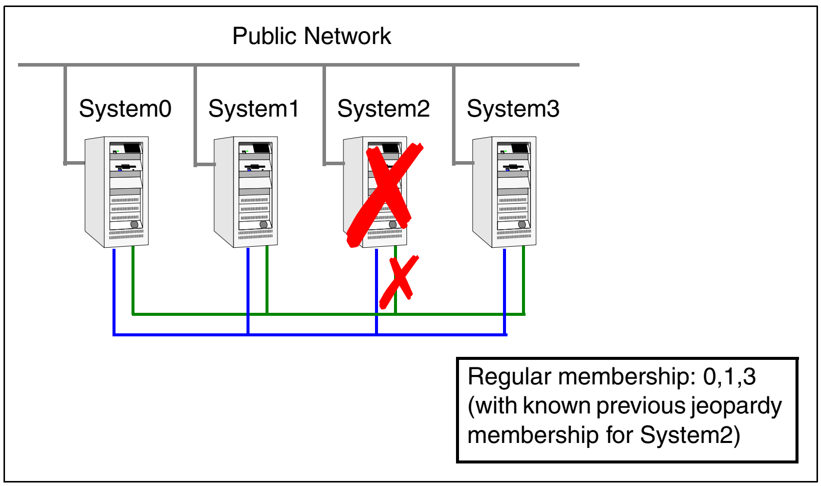

Systems 0, 1, and 3 recognize that System2 has faulted. The cluster is reformed. Systems 0, 1, and 3 are in a regular membership. When System2 went into jeopardy membership, service groups running on System2 were autodisabled. Even though the system is now completely failed, no other system can assume ownership of these service groups unless the system administrator manually clears the AutoDisabled flag on the service groups that were running on System2.

However, after the flag is cleared, these service groups can be manually brought online on other systems in the cluster.

All high priority cluster interconnect links fail

In this example, all high priority links to System2 fail. This can occur two ways:

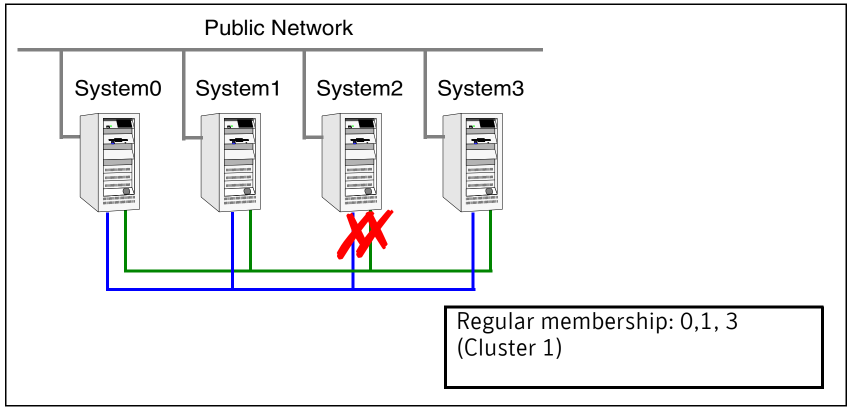

- Both links to System2 fail at the same time
  System2 was never in jeopardy membership. Without a low priority link, the cluster splits into two subclusters, where System0, 1 and 3 are in one subcluster, and System2 is in another. This is a split brain scenario.
- Both links to System2 fail at different times
  System2 was in a jeopardy membership when the second link failed, and therefore the service groups that were online on System2 were autodisabled. No other system can online these service groups without administrator intervention.
  Systems 0, 1 and 3 form a mini-cluster. System2 forms another single-system mini-cluster. All service groups that were present on systems 0, 1 and 3 are autodisabled on System2.

Four system cluster with low priority link

Consider a four-system cluster that has two private cluster interconnect heartbeat links, and one public low priority link.

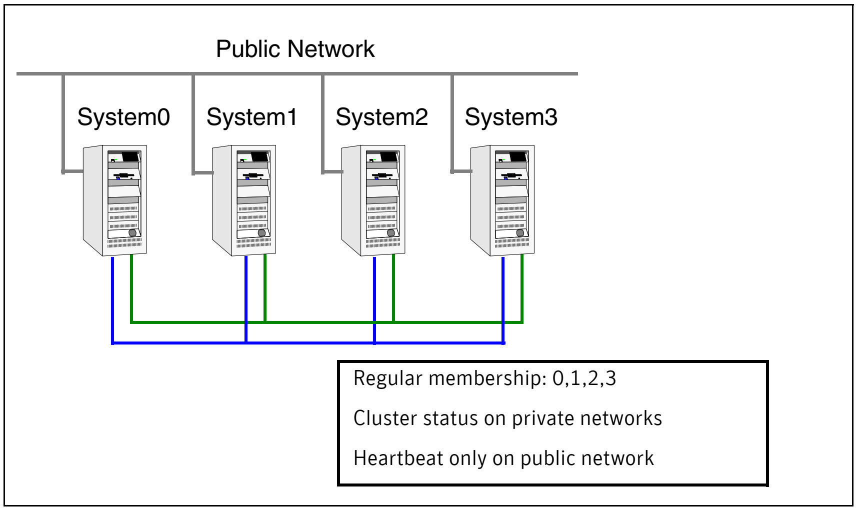

Cluster interconnect link failure

In this example, a link to System2 fails, leaving System2 with one cluster interconnect link and the low priority link remaining.

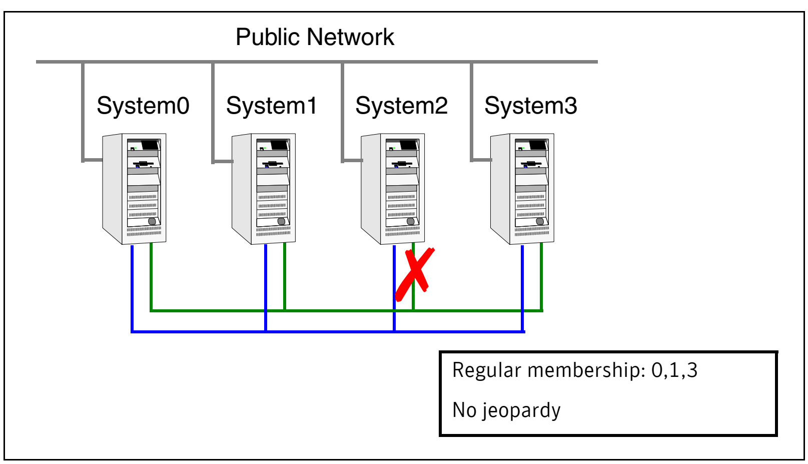

Other systems send all cluster status traffic to System2 over the remaining private link and use both private links for traffic between themselves. The low priority link continues carrying the heartbeat signal only. No jeopardy condition is in effect because two links remain to determine system failure.

Cluster interconnect link failure followed by system failure

In this example, the link to System2 fails. Because there is a low priority heartbeat link, System2 is not put in the jeopardy membership. Subsequently, System2 fails due to a power fault.

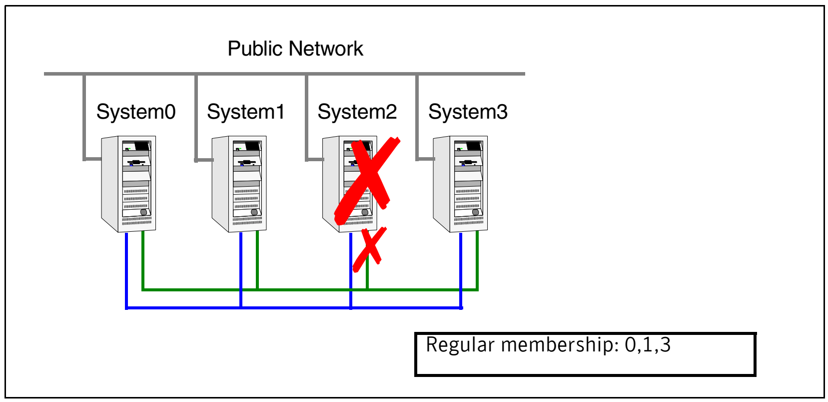

Systems 0, 1, and 3 recognize that System2 has faulted. The cluster is reformed. Systems 0, 1, and 3 are in a regular membership. The service groups on System2 that are configured for failover on system fault are attempted to be brought online on another target system, if one exists.

All high priority cluster interconnect links fail

In this example, both high priority cluster interconnect links to System2 fail, leaving System2 with only the low priority link remaining.

Cluster status communication is now routed over the low priority link to System2. System2 is placed in a jeopardy membership. The service groups on

System2 are autodisabled, and the service group attribute AutoFailOver is set to 0, meaning the service group will not fail over on a system fault.

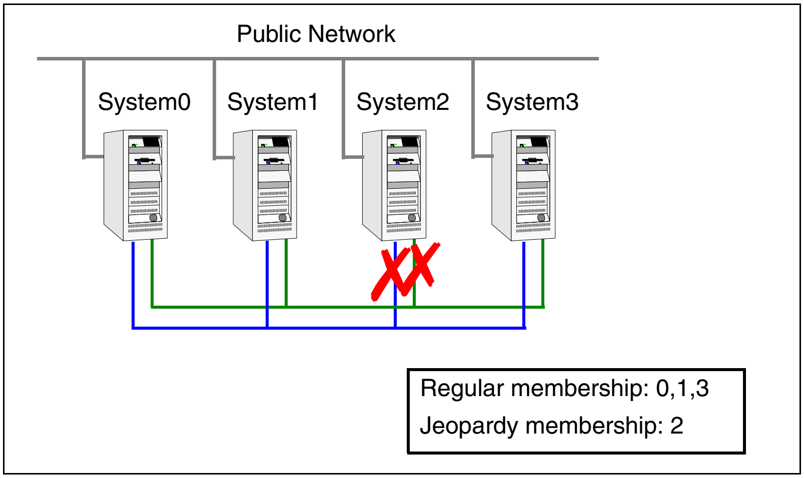

When a cluster interconnect link is re-established, all cluster status communications revert back to the cluster interconnect and the low priority link returns to sending heartbeat signal only. At this point, System2 is placed back in regular cluster membership.

Summary of best practices for cluster communications

The following are the recommended best practices for cluster communications to best support proper cluster membership and data protection.

- Properly seed the cluster by requiring all systems, and not just a subset of systems, to be present in the GAB membership before the cluster will automatically seed.
  If every system is not present, manual intervention by the administrator must eliminate the possibility of a split brain condition before manually seeding the cluster.
- Configure multiple independent communication network links between cluster systems.
  Networks should not have a single point of failure, such as a shared hub or ethernet card.
- Low-priority LLT links in clusters with or without I/O fencing is recommended. In clusters without I/O fencing, this is critical.

Note: An exception to this is if the cluster uses fencing along with Cluster File Systems (CFS) or Oracle Real Application Clusters (RAC).

The reason for this is that low priority links are usually shared public network links. In the case where the main cluster interconnects fail, and the low priority link was the only remaining link, large amounts of data would be moved to the low priority link. This would potentially slow down the public network to unacceptable performance. Without a low priority link configured, membership arbitration would go into effect in this case, and some systems may be taken down, but the remaining systems would continue to run without impact to the public network.

It is not recommended to have a cluster with CFS or RAC without I/O fencing configured.

- Disable the console-abort sequence

  Most UNIX systems provide a console-abort sequence that enables the administrator to halt and continue the processor. Continuing operations after the processor has stopped may corrupt data and is therefore unsupported by VCS.

  When a system is halted with the abort sequence, it stops producing heartbeats. The other systems in the cluster consider the system failed and take over its services. If the system is later enabled with another console sequence, it continues writing to shared storage as before, even though its applications have been restarted on other systems.

  Symantec recommends disabling the console-abort sequence or creating an alias to force the go command to perform a restart on systems not running I/O fencing.

- Select the smallest possible LUNs for use as coordinator disks. No more than three coordinator disks are needed in any configuration.

- Do not reconnect the cluster interconnect after a network partition without shutting down one side of the split cluster.

  A common example of this happens during testing, where the administrator may disconnect the cluster interconnect and create a network partition. Depending on when the interconnect cables are reconnected, unexpected behavior can occur.

Chapter 13

Controlling VCS behavior

- About VCS behavior on resource faults
- Controlling VCS behavior at the service group level
- Controlling VCS behavior at the resource level
- Changing agent file paths and binaries
- Service group workload management
- Sample configurations depicting workload management

About VCS behavior on resource faults

VCS considers a resource faulted in the following situations:

- When the resource state changes unexpectedly. For example, an online resource going offline.
- When a required state change does not occur. For example, a resource failing to go online or offline when commanded to do so.

In many situations, VCS agents take predefined actions to correct the issue before reporting resource failure to the engine. For example, the agent may try to bring a resource online several times before declaring a fault.

When a resource faults, VCS takes automated actions to “clean up the faulted resource. The Clean function makes sure the resource is completely shut down before bringing it online on another node. This prevents concurrency violations.

When a resource faults, VCS takes all resources dependent on the faulted resource offline. The fault is thus propagated in the service group

About critical and non-critical resources

The Critical attribute for a resource defines whether a service group fails over when the resource faults. If a resource is configured as non-critical (by setting the Critical attribute to 0) and no resources depending on the failed resource are critical, the service group will not fail over. VCS takes the failed resource offline and updates the group status to ONLINE|PARTIAL. The attribute also determines whether a service group tries to come online on another node if, during the group’s online process, a resource fails to come online.

VCS behavior diagrams

This section describes the default functionality of VCS when resources fault. The illustration displays the symbols used in this section.

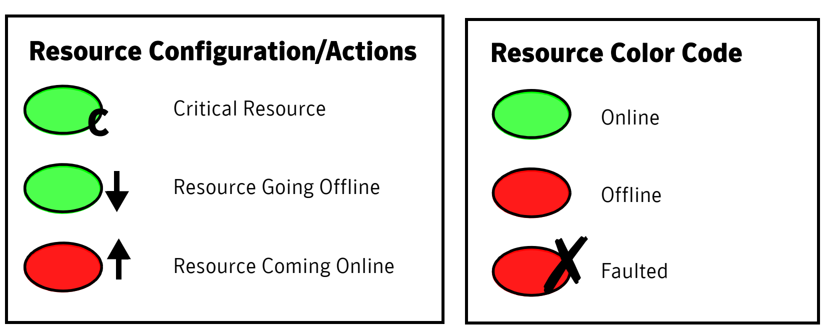

Scenario: Resource with critical parent faults

The service group in the following example has five resources, of which resource R1 is configured as a critical resource.

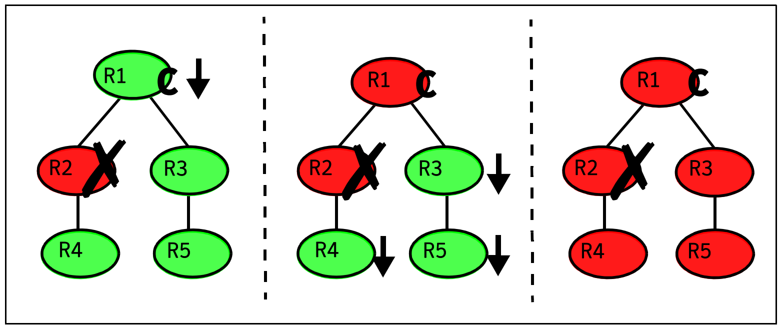

When resource R2 faults, the fault is propagated up the dependency tree to resource R1. When the critical resource R1 goes offline, VCS must fault the service group and fail it over elsewhere in the cluster. VCS takes other resources in the service group offline in the order of their dependencies. After taking resources R3, R4, and R5 offline, VCS fails over the service group to another node.

Scenario: Resource with non-critical parent faults

The service group in the following example does not have any critical resources.

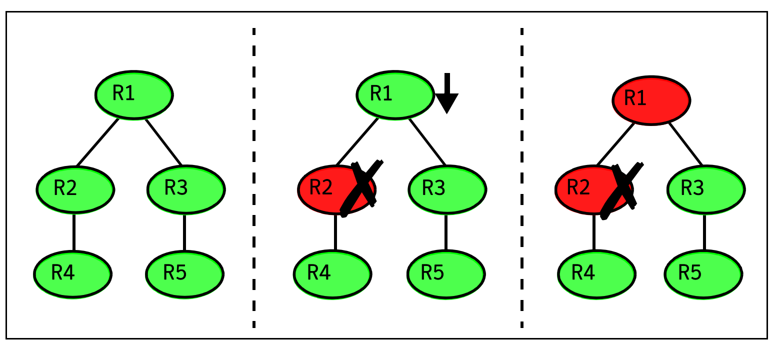

When resource R2 faults, the engine propagates the failure up the dependency tree. Neither resource R1 nor resource R2 are critical, so the fault does not result in the tree going offline or in service group failover.

Scenario: Resource with critical parent fails to come online

In the following example, when a command is issued to bring the service group online, resource R2 fails to come online.

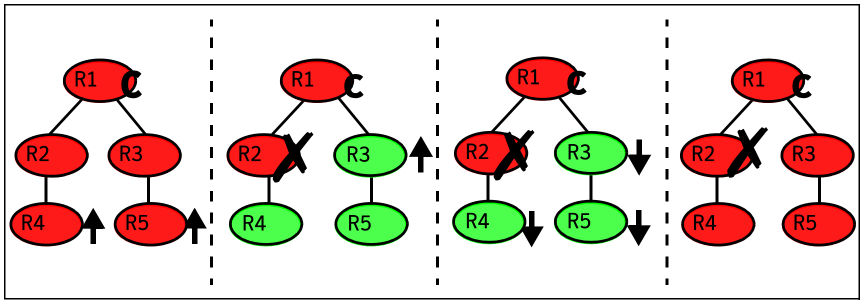

VCS calls the Clean function for resource R2 and propagates the fault up the dependency tree. Resource R1 is set to critical, so the service group is taken offline and failed over to another node in the cluster.

Controlling VCS behavior at the service group level

This section describes how you can configure service group attributes to modify VCS behavior in response to resource faults.

About the AutoRestart attribute

If a persistent resource on a service group (GROUP_1) faults, VCS fails the service group over to another system if the following conditions are met:

- The AutoFailOver attribute is set.
- Another system in the cluster exists to which GROUP_1 can fail over.

If neither of these conditions is met, GROUP_1 remains offline and faulted, even after the faulted resource becomes online.

Setting the AutoRestart attribute enables a service group to be brought back online without manual intervention. Or if GROUP_1 could no failover targets are available, setting the AutoRestart attribute enables VCS to bring the group back online on the first available system after the group's faulted resource came online on that system.

For example, NIC is a persistent resource. In some cases, when a system boots and VCS starts, VCS probes all resources on the system. When VCS probes the NIC resource, the resource may not be online because the networking is not up and fully operational. In such situations, VCS marks the NIC resource as faulted, and does not bring the service group online. However, when the NIC resource becomes online and if AutoRestart is enabled, the service group is brought online.

Controlling failover on service group or system faults

The AutoFailOver attribute configures service group behavior in response to service group and system faults.

- If the AutoFailOver attribute is set to 1, the service group fails over when a system or a service group faults, provided a suitable system exists for failover.
- If the AutoFailOver attribute is set to 0, the service group does not fail over when a system or service group faults. If a fault occurs in a service group, the group is taken offline, depending on whether any of its resources are configured as critical. If a system faults, the service group is not failed over to another system.

Defining failover policies

The service group attribute FailOverPolicy governs how VCS calculates the target system for failover. There are three possible values for FailOverPolicy:

Priority	VCS selects the system with the lowest priority as the failover target. The Priority failover policy is ideal for simple two-node clusters or small clusters with few service groups. Priority is set in the SystemList attribute implicitly via ordering, such as SystemList = {SystemA, SystemB} or explicitly, such as SystemList = {SystemA=0, SystemB=1}. Priority is the default behavior.
RoundRobin	VCS selects the system running the fewest service groups as the failover target. This policy is ideal for large clusters running many service groups with similar server load characteristics (for example, similar databases or applications)
Load	The Load failover policy comprises the following components: System capacity and service group load, represented by the attributes Capacity and Load respectively. System limits and service group prerequisites, represented by the attributes Limits and Prerequisites, respectively.

About system zones

The SystemZones attribute enables you to create a subset of systems to use in an initial failover decision. This feature allows fine-tuning of application failover decisions, and yet retains the flexibility to fail over anywhere in the cluster.

If the attribute is configured, a service group tries to stay within its zone before choosing a host in another zone. For example, in a three-tier application infrastructure with Web, application, and database servers, you could create two system zones: one each for the application and the database. In the event of a failover, a service group in the application zone will try to fail over to another node within the zone. If no nodes are available in the application zone, the group will fail over to the database zone, based on the configured load and limits.

In this configuration, excess capacity and limits on the database backend are kept in reserve to handle the larger load of a database failover. The application servers handle the load of service groups in the application zone. During a cascading failure, the excess capacity in the cluster is available to all service groups.

About load-based autostart

VCS provides a method to determine where a service group comes online when the cluster starts. Setting the AutoStartPolicy to Load instructs the VCS engine, HAD, to determine the best system on which to start the groups. VCS places service groups in an AutoStart queue for load-based startup as soon as the groups probe all running systems. VCS creates a subset of systems that meet all prerequisites and then chooses the system with the highest AvailableCapacity.

Set AutoStartPolicy = Load and configure the SystemZones attribute to establish a list of preferred systems on which to initially run a group.

Freezing service groups

Freezing a service group prevents VCS from taking any action when the service group or a system faults. Freezing a service group prevents dependent resources from going offline when a resource faults. It also prevents the Clean function from being called on a resource fault.

You can freeze a service group when performing operations on its resources from outside VCS control. This prevents VCS from taking actions on resources while your operations are on. For example, freeze a database group when using database controls to stop and start a database.

Controlling Clean behavior on resource faults

The ManageFaults attribute specifies whether VCS calls the Clean entry point when a resource faults. ManageFaults is a service group attribute; you can configure each service group to operate as desired.

- If the ManageFaults attribute is set to ALL, VCS calls the Clean entry point when a resource faults.
- If the ManageFaults attribute is set to NONE, VCS takes no action on a resource fault; it "hangs the service group until administrative action can be taken. VCS marks the resource state as ADMIN_WAIT and does not fail over the service group until the resource fault is removed and the ADMIN_WAIT state is cleared.

  VCS calls the resadminwait trigger when a resource enters the ADMIN_WAIT state due to a resource fault if the ManageFaults attribute is set to NONE. You can customize this trigger to provide notification about the fault.

  See "resadminwait event trigger" on page 533.

  When ManageFaults is set to none and one of the following events occur, the resource enters the admin_wait state:

Event	Resource state
The offline entry point did not complete within the expected time.	ONLINE|ADMIN_WAIT
The offline entry point was ineffective.	ONLINE|ADMIN_WAIT
The online entry point did not complete within the expected time.	OFFLINE|ADMIN_WAIT
The online entry point was ineffective.	OFFLINE|ADMIN_WAIT
The resource was taken offline unexpectedly.	OFFLINE|ADMIN_WAIT
For the online resource the monitor entry point consistently failed to complete within the expected time.	ONLINE|MONITOR_TIMEDOUT|ADMIN_WAIT

Clearing resources in the ADMIN_WAIT state

When VCS sets a resource in the ADMIN_WAIT state, it invokes the resadminwait trigger according to the reason the resource entered the state.

See "resadminwait event trigger" on page 533.

To clear a resource

1. Take the necessary actions outside VCS to bring all resources into the required state.
2. Verify that resources are in the required state by issuing the command:

   hagrp -clearadminwait *group* -sys *system*

   This command clears the ADMIN_WAIT state for all resources. If VCS continues to detect resources that are not in the required state, it resets the resources to the ADMIN_WAIT state.
3. If resources continue in the ADMIN_WAIT state, repeat step 1 and step 2, *or* issue the following command to stop VCS from setting the resource to the ADMIN_WAIT state:

   hagrp -clearadminwait -fault *group* -sys *system*

   This command has the following results:

   - If the resadminwait trigger was called for reasons 0 or 1, the resource state is set as ONLINE|UNABLE_TO_OFFLINE.
   - If the resadminwait trigger was called for reasons 2, 3, or 4, the resource state is set as FAULTED. Please note that when resources are set as FAULTED for these reasons, the clean entry point is not called. Verify that resources in ADMIN-WAIT are in clean, OFFLINE state prior to invoking this command.

   When a service group has a resource in the ADMIN_WAIT state, the following service group operations cannot be performed on the resource: online, offline, switch, and flush. Also, you cannot use the hastop command when resources are in the ADMIN_WAIT state. When this occurs, you must issue the hastop command with -force option only.

Controlling fault propagation

The FaultPropagation attribute defines whether a resource fault is propagated up the resource dependency tree. It also defines whether a resource fault causes a service group failover.

- If the FaultPropagation attribute is set to 1 (default), a resource fault is propagated up the dependency tree. If a resource in the path is critical, the service group is taken offline and failed over, provided the AutoFailOver attribute is set to 1.
- If the FaultPropagation is set to 0, resource faults are contained at the resource level. VCS does not take the dependency tree offline, thus preventing failover. If the resources in the service group remain online, the service group remains in the PARTIAL|FAULTED state. If all resources are offline or faulted, the service group remains in the OFFLINE| FAULTED state.

When a resource faults, VCS fires the resfault trigger and sends an SNMP trap. The trigger is called on the system where the resource faulted and includes the name of the faulted resource.

See"resfault event trigger" on page 535.

Customized behavior diagrams

The illustrations in this section depict how the ManageFaults and FaultPropagation attributes change VCS behavior when handling resource faults The following illustration depicts the legends used in the section.

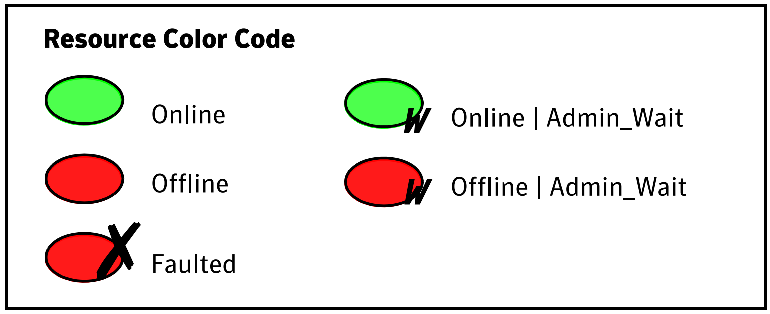

Scenario: Resource with a critical parent and ManageFaults=NONE

The service group in the following example has five resources. The ManageFaults attribute is set to NONE for resource R2.

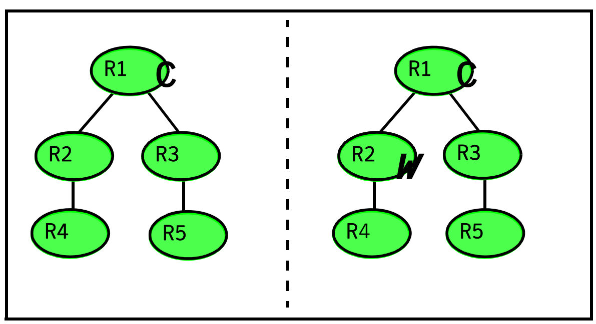

If resource R2 fails, the resource is marked as ONLINE|ADMIN_WAIT. The Clean entry point is not called for the resource. VCS does not take any other resource offline.

Scenario: Resource with a critical parent and FaultPropagation=0

In the following example, the FaultPropagation attribute is set to 0.

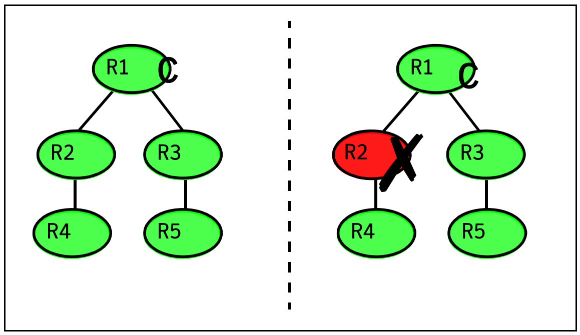

When resource R2 faults, the Clean entry point is called and the resource is marked as faulted. The fault is not propagated up the tree, and the group is not taken offline.

Controlling VCS behavior at the resource level

This section describes how you can control VCS behavior at the resource level. Note that a resource is not considered faulted until the agent framework declares the fault to the VCS engine.

About resource type attributes that control resource behavior

The following attributes affect how the VCS agent framework reacts to problems with individual resources before informing the fault to the VCS engine.

About the RestartLimit attribute

The RestartLimit attribute defines whether VCS attempts to restart a failed resource before informing the engine of the fault.

If the RestartLimit attribute is set to a non-zero value, the agent attempts to restart the resource before declaring the resource as faulted. When restarting a failed resource, the agent framework calls the Clean entry point before calling the Online entry point. However, setting the ManageFaults attribute to NONE prevents the Clean entry point from being called and prevents the Online entry point from being retried.

About the OnlineRetryLimit attribute

The OnlineRetryLimit attribute specifies the number of times the Online entry point is retried if the initial attempt to bring a resource online is unsuccessful.

When the OnlineRetryLimit set to a non-zero value, the agent framework calls the Clean entry point before rerunning the Online entry point. Setting the ManageFaults attribute to NONE prevents the Clean entry point from being called and also prevents the Online operation from being retried.

About the ConfInterval attribute

The ConfInterval attribute defines how long a resource must remain online without encountering problems before previous problem counters are cleared. The attribute controls when VCS clears the RestartCount, ToleranceCount and CurrentMonitorTimeoutCount values.

About the ToleranceLimit attribute

The ToleranceLimit attribute defines the number of times the Monitor routine should return an offline status before declaring a resource offline. This attribute is typically used when a resource is busy and appears to be offline. Setting the attribute to a non-zero value instructs VCS to allow multiple failing monitor cycles with the expectation that the resource will eventually respond. Setting a non-zero ToleranceLimit also extends the time required to respond to an actual fault.

About the FaultOnMonitorTimeouts attribute

The FaultOnMonitorTimeouts attribute defines whether VCS interprets a Monitor entry point timeout as a resource fault.

If the attribute is set to 0, VCS does not treat Monitor timeouts as a resource faults. If the attribute is set to 1, VCS interprets the timeout as a resource fault and the agent calls the Clean entry point to shut the resource down.

By default, the FaultOnMonitorTimeouts attribute is set to 4. This means that the Monitor entry point must time out four times in a row before the resource is marked faulted.

How VCS handles resource faults

This section describes the process VCS uses to determine the course of action when a resource faults.

VCS behavior when an online resource faults

In the following example, a resource in an online state is reported as being offline without being commanded by the agent to go offline.

- VCS first verifies the Monitor routine completes successfully in the required time. If it does, VCS examines the exit code returned by the Monitor routine. If the Monitor routine does not complete in the required time, VCS looks at the FaultOnMonitorTimeouts (FOMT) attribute.
- If FOMT=0, the resource will not fault when the Monitor routine times out. VCS considers the resource online and monitors the resource periodically, depending on the monitor interval.
  If FOMT=1 or more, VCS compares the CurrentMonitorTimeoutCount (CMTC) with the FOMT value. If the monitor timeout count is not used up, CMTC is incremented and VCS monitors the resource in the next cycle.
- If FOMT= CMTC, this means that the available monitor timeout count is exhausted and VCS must now take corrective action.
- If the ManageFaults attribute is set to NONE, VCS marks the resource as ONLINE|ADMIN_WAIT and fires the resadminwait trigger. If the ManageFaults attribute is set to ALL, the resource enters a GOING OFFLINE WAIT state. VCS invokes the Clean entry point with the reason *Monitor Hung*.

- If the Clean entry point is successful (that is, Clean exit code = 0), VCS examines the value of the RestartLimit attribute. If Clean fails (exit code = 1), the resource remains online with the state UNABLE TO OFFLINE. VCS fires the resnotoff trigger and monitors the resource again.
- If the Monitor routine does not time out, it returns the status of the resource as being online or offline.
- If the ToleranceLimit (TL) attribute is set to a non-zero value, the Monitor cycle returns offline (exit code = 100) for a number of times specified by the ToleranceLimit and increments the ToleranceCount (TC). When the ToleranceCount equals the ToleranceLimit (TC = TL), the agent declares the resource as faulted.
- If the Monitor routine returns online (exit code = 110) during a monitor cycle, the agent takes no further action. The ToleranceCount attribute is reset to 0 when the resource is online for a period of time specified by the ConfInterval attribute.

  If the resource is detected as being offline a number of times specified by the ToleranceLimit before the ToleranceCount is reset (TC = TL), the resource is considered failed.
- After the agent determines the resource is not online, VCS checks the Frozen attribute for the service group. If the service group is frozen, VCS declares the resource faulted and calls the resfault trigger. No further action is taken.
- If the service group is not frozen, VCS checks the ManageFaults attribute. If ManageFaults=NONE, VCS marks the resource state as ONLINE|ADMIN_WAIT and calls the resadminwait trigger. If ManageFaults=ALL, VCS calls the Clean entry point with the CleanReason set to Unexpected Offline.
- If the Clean entry point fails (exit code = 1) the resource remains online with the state UNABLE TO OFFLINE. VCS fires the resnotoff trigger and monitors the resource again. The resource enters a cycle of alternating Monitor and Clean entry points until the Clean entry point succeeds or a user intervenes.
- If the Clean entry point is successful, VCS examines the value of the RestartLimit (RL) attribute. If the attribute is set to a non-zero value, VCS increments the RestartCount (RC) attribute and invokes the Online entry point. This continues till the value of the RestartLimit equals that of the RestartCount. At this point, VCS attempts to monitor the resource.

If the Monitor returns an online status, VCS considers the resource online and resumes periodic monitoring. If the monitor returns an offline status, the resource is faulted and VCS takes actions based on the service group configuration.

VCS behavior when a resource fails to come online

In the following example, the agent framework invokes the Online entry point for an offline resource. The resource state changes to WAITING TO ONLINE.

- If the Online entry point times out, VCS examines the value of the ManageFaults attribute.
- If ManageFaults is set to NONE, the resource state changes to OFFLINE|ADMIN_WAIT.
  If ManageFaults is set to ALL, VCS calls the Clean entry point with the CleanReason set to Online Hung.
- If the Online entry point does not time out, VCS invokes the Monitor entry point. The Monitor routine returns an exit code of 110 if the resource is online. Otherwise, the Monitor routine returns an exit code of 100.
- VCS examines the value of the OnlineWaitLimit (OWL) attribute. This attribute defines how many monitor cycles can return an offline status before the agent framework declares the resource faulted. Each successive Monitor cycle increments the OnlineWaitCount (OWC) attribute. When OWL= OWC (or if OWL= 0), VCS determines the resource has faulted.
- VCS then examines the value of the ManageFaults attribute. If the ManageFaults is set to NONE, the resource state changes to OFFLINE|ADMIN_WAIT.
  If the ManageFaults is set to ALL, VCS calls the Clean entry point with the CleanReason set to Online Ineffective.
- If the Clean entry point is not successful (exit code = 1), the agent monitors the resource. It determines the resource is offline, and calls the Clean entry point with the Clean Reason set to Online Ineffective. This cycle continues till the Clean entry point is successful, after which VCS resets the OnlineWaitCount value.
- If the OnlineRetryLimit (ORL) is set to a non-zero value, VCS increments the OnlineRetryCount (ORC) and invokes the Online entry point. This starts the

cycle all over again. If ORL = ORC, or if ORL = 0, VCS assumes that the Online operation has failed and declares the resource as faulted.

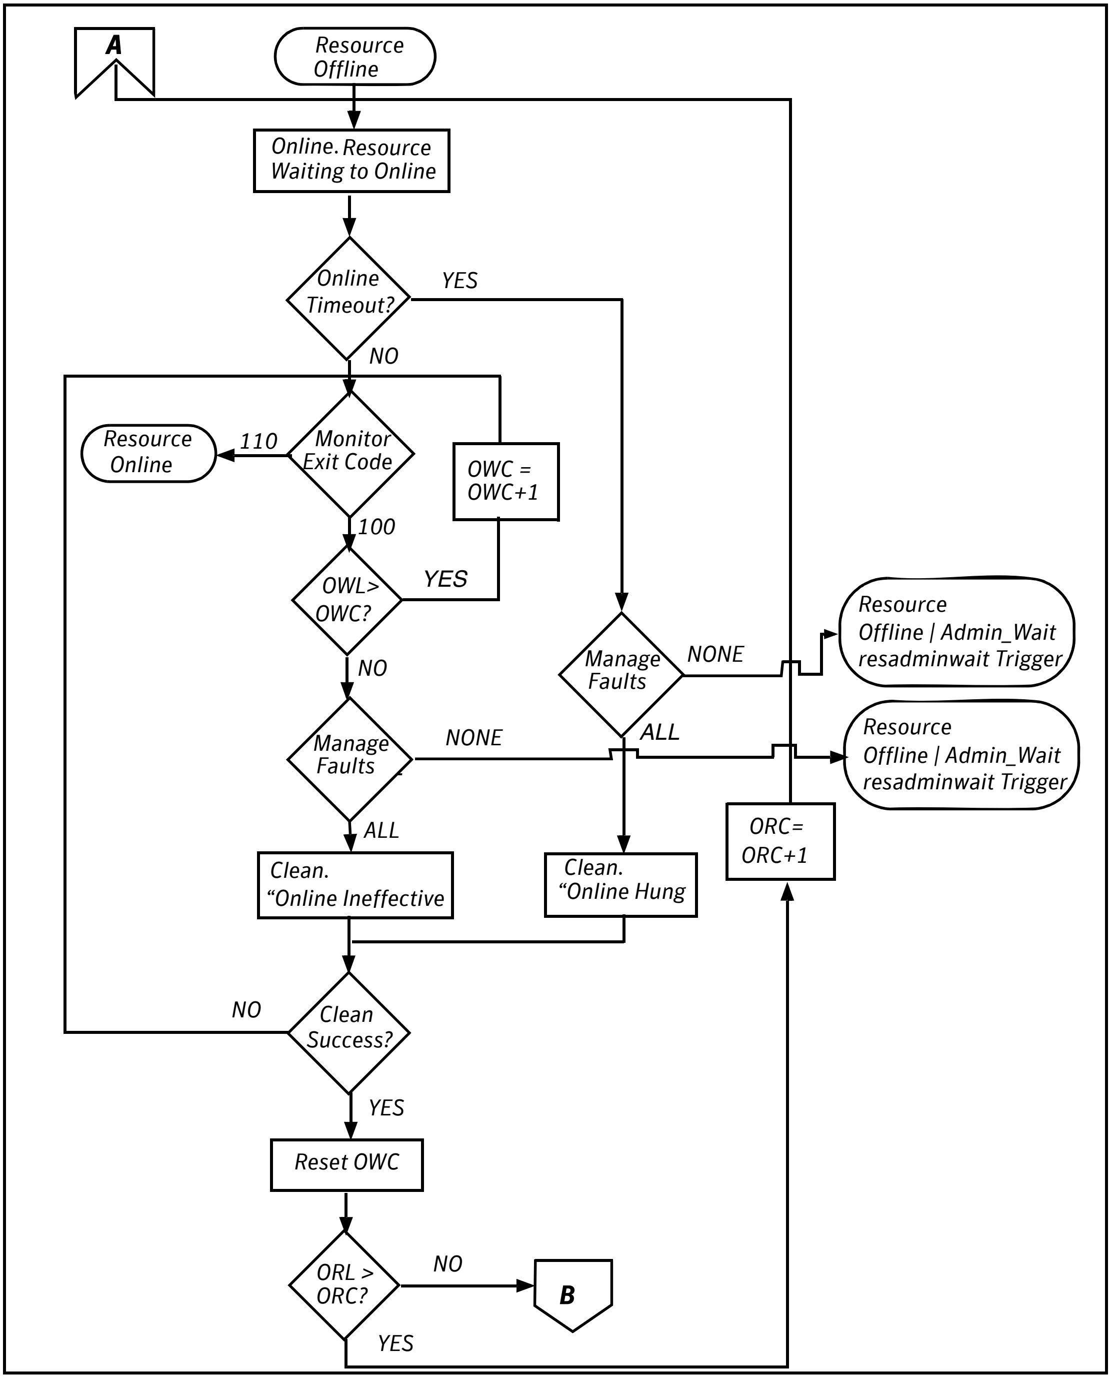
A
Resource Offline
Online. Resource Waiting to Online
Online Timeout?
YES
NO
Resource Online
110
Monitor Exit Code
OWC = OWC+1
100
OWL> OWC?
YES
NO
Manage Faults
NONE
Resource Offline | Admin_Wait resadminwait Trigger
Manage Faults
NONE
ALL
Resource Offline | Admin_Wait resadminwait Trigger
ALL
Clean. “Online Ineffective
Clean. “Online Hung
ORC= ORC+1
NO
Clean Success?
YES
Reset OWC
ORL > ORC?
NO
B
YES

VCS behavior after a resource is declared faulted

After a resource is declared faulted, VCS fires the resfault trigger and examines the value of the FaultPropagation attribute.

- If FaultPropagation is set to 0, VCS does not take other resources offline, and changes the group state to OFFLINE|FAULTED or PARTIAL|FAULTED. The service group does not fail over.
  If FaultPropagation is set to 1, VCS takes all resources in the dependent path of the faulted resource offline, up to the top of the tree.
- VCS then examines if any resource in the dependent path is critical. If no resources are critical, the service group is left in its OFFLINE|FAULTED or PARTIAL|FAULTED state. If a resource in the path is critical, VCS takes the all resources in the service group offline in preparation of a failover.
- If the AutoFailOver attribute is set to 0, the service group is not failed over; it remains in a faulted state. If AutoFailOver is set to 1, VCS examines if any systems in the service group's SystemList are possible candidates for failover. If no suitable systems exist, the group remains faulted and VCS calls the nofailover trigger. If eligible systems are available, VCS examines the FailOverPolicy to determine the most suitable system to which to fail over the service group.

- If FailOverPolicy is set to Load, a NoFailover situation may occur because of restrictions placed on service groups and systems by Service Group Workload Management.

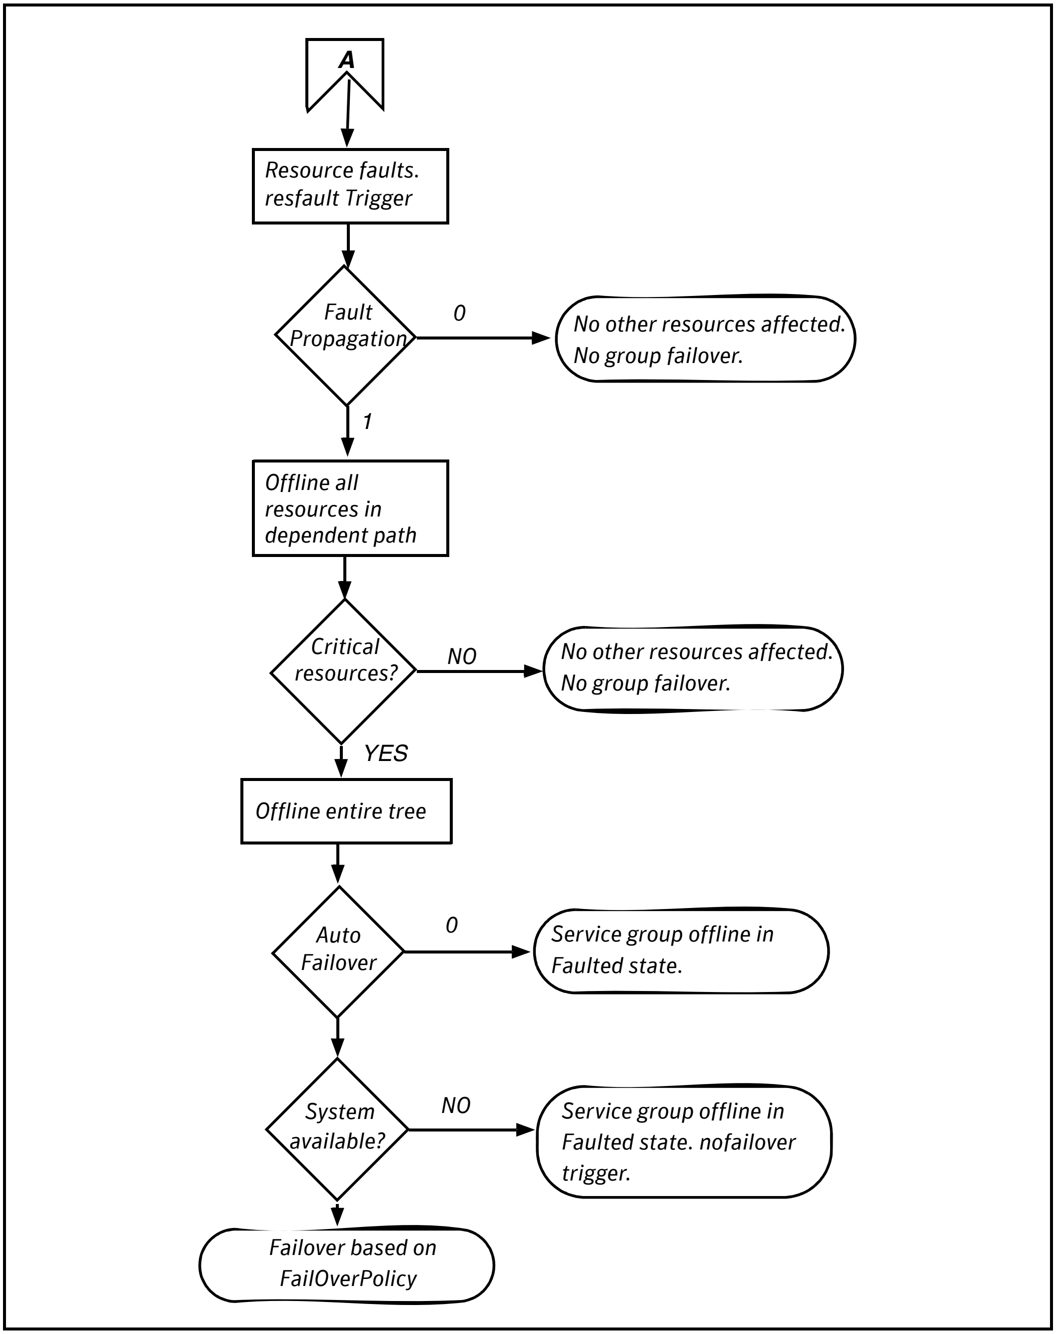
A
Resource faults. resfault Trigger
Fault Propagation
0
No other resources affected. No group failover.
1
Offline all resources in dependent path
Critical resources?
NO
No other resources affected. No group failover.
YES
Offline entire tree
Auto Failover
0
Service group offline in Faulted state.
System available?
NO
Service group offline in Faulted state. nofailover trigger.
Failover based on FailOverPolicy

Disabling resources

Disabling a resource means that the resource is no longer monitored by a VCS agent, and that the resource cannot be brought online or taken offline. The agent starts monitoring the resource after the resource is enabled. The resource attribute Enabled determines whether a resource is enabled or disabled. A persistent resource can be disabled when all its parents are offline. A non-persistent resource can be disabled when the resource is in an OFFLINE state.

When to disable a resource

Typically, resources are disabled when one or more resources in the service group encounter problems and disabling the resource is required to keep the service group online or to bring it online.

Note: Disabling a resource is not an option when the entire service group requires disabling. In that case, set the service group attribute Enabled to 0.

To disable a resource

To disable the resource when VCS is running:

```
hares -modify resource_name Enabled 0
```

To have the resource disabled initially when VCS is started, set the resource's Enabled attribute to 0 in main.cf.

Limitations of disabling resources

When VCS is running, there are certain prerequisites to be met before the resource is disabled successfully.

- An online non-persistent resource cannot be disabled. It must be in a clean OFFLINE state. (The state must be OFFLINE and IState must be NOT WAITING.)
- If it is a persistent resource and the state is ONLINE on some of the systems, all dependent resources (parents) must be in clean OFFLINE state. (The state must be OFFLINE and IState must be NOT WAITING)

Therefore, before disabling the resource you may be required to take it offline (if it is non-persistent) and take other resources offline in the service group.

Additional considerations for disabling resources

- When a group containing disabled resources is brought online, the online transaction is not propagated to the disabled resources. Children of the disabled resource are brought online by VCS only if they are required by another enabled resource.
- You can bring children of disabled resources online if necessary.
- When a group containing disabled resources is taken offline, the offline transaction is propagated to the disabled resources.

The following figures show how a service group containing disabled resources is brought online.

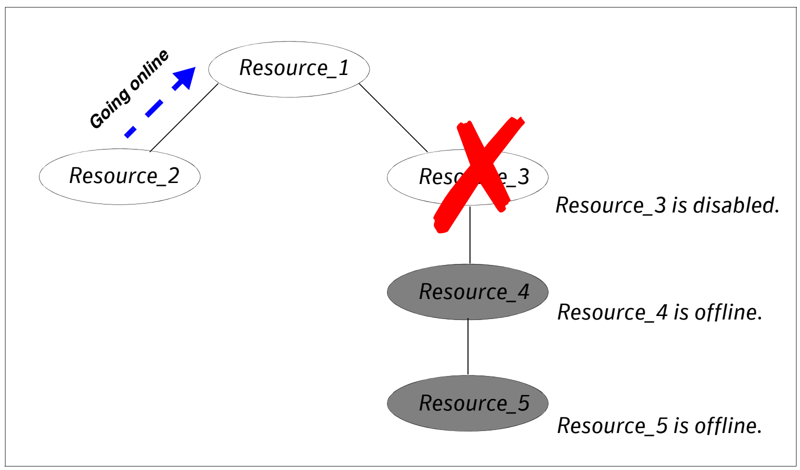

In the figure above, Resource_3 is disabled. When the service group is brought online, the only resources brought online by VCS are Resource_1 and Resource_2 (Resource_2 is brought online first) because VCS recognizes Resource_3 is disabled. In accordance with online logic, the transaction is not propagated to the disabled resource.

In the figure below, Resource_2 is disabled. When the service group is brought online, resources 1, 3, 4 are also brought online (Resource_4 is brought online first). Note Resource_3, the child of the disabled resource, is brought online because Resource_1 is enabled and is dependent on it.

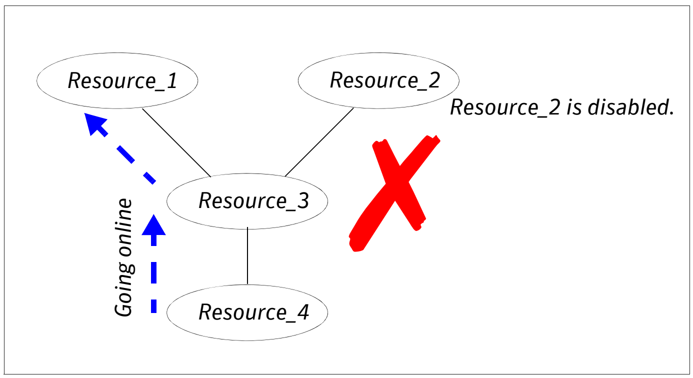

How disabled resources affect group states

When a service group is brought online containing non-persistent, disabled resources whose AutoStart attributes are set to 1, the group state is PARTIAL, even though enabled resources with Autostart=1 are online. This is because the disabled resource is considered for the group state.

To have the group in the ONLINE state when enabled resources with AutoStart set to 1 are in ONLINE state, set the AutoStart attribute to 0 for the disabled, non-persistent resources.

Changing agent file paths and binaries

VCS runs agent binaries from the path `%VCS_HOME%\bin\agent_name\`.

You can instruct VCS to run a different set of agent binaries or scripts by specifying values for the AgentFile and AgentDirectory attributes.

- AgentFile— Specify a value for this attribute if the name of the agent binary is not the same as that of the resource type.
  For example, if the resource type is NetBackup and the agent binary is called NBU.dll, set the AgentFile attribute to NBU.dll.
- AgentDirectory— Specify a value for this attribute if the agent is not installed at the default location.
  When you specify the agent directory, VCS looks for the agent file (*AgentName*Agent) in the agent directory. If the agent file name does not conform to the *AgentName*Agent convention, configure the AgentFile attribute.
  For example, if the NetBackup agent is installed at *C:\Program Files\VERITAS\NetBackup*, specify this path as the attribute value.

To change the path of an agent

◆ Before configuring a resource for the agent, add AgentFile and AgentDirectory as static attributes to the agent's resource type.

```
haattr -add -static resource_type AgentFile "binary_name.dll"
haattr -add -static resource_type AgentDirectory "C:\Program
    Files\agent_path"
```

Service group workload management

Workload management is a load-balancing mechanism that determines which system hosts an application during startup, or after an application or server fault.

Service Group Workload Management provides tools for making intelligent decisions about startup and failover locations, based on system capacity and resource availability.

Enabling service group workload management

The service group attribute FailOverPolicy governs how VCS calculates the target system for failover. Set FailOverPolicy to Load to enable service group workload management.

See "Controlling VCS behavior at the resource level" on page 459

About system capacity and service group load

The Load and Capacity construct allows the administrator to define a fixed amount of resources a server provides (Capacity), and a fixed amount of resources a specific service group is expected to utilize (Load).

The system attribute Capacity sets a fixed load-handling capacity for servers. Define this attribute based on system requirements.

The service group attribute Load sets a fixed demand for service groups. Define this attribute based on application requirements.

When a service group is brought online, its load is subtracted from the system's capacity to determine available capacity. VCS maintains this info in the attribute AvailableCapacity.

When a failover occurs, VCS determines which system has the highest available capacity and starts the service group on that system. During a failover involving multiple service groups, VCS makes failover decisions serially to facilitate a proper load-based choice.

System capacity is a *soft* restriction; in some situations, value of the Capacity attribute could be less than zero. During some operations, including cascading failures, the value of the AvailableCapacity attribute could be negative.

Static load versus dynamic load

Dynamic load is an integral component of the Service Group Workload Management framework. Typically, HAD sets remaining capacity with the function:

AvailableCapacity = Capacity - (sum of Load values of all online service groups)

If the DynamicLoad attribute is defined, its value overrides the calculated Load values with the function:

AvailableCapacity = Capacity - DynamicLoad

This enables better control of system loading values than estimated service group loading (static load). However, this requires setting up and maintaining a load estimation package outside VCS. It also requires modifying the configuration file main.cf manually.

Note that the DynamicLoad (specified with `hasys -load`) is subtracted from the Capacity as an integer and not a percentage value. For example, if a system's capacity is 200 and the load estimation package determines the server is 80 percent loaded, it must inform VCS that the DynamicLoad value is 160 (not 80).

About overload warning

Overload warning provides the notification component of the Load policy. When a server sustains the preset load level (set by the attribute LoadWarningLevel) for a preset time (set by the attribute LoadTimeThreshold), VCS invokes the loadwarning trigger.

See "Using event triggers" on page 528

See "System attributes" on page 728.

The loadwarning trigger is a user-defined script or application designed to carry out specific actions. It is invoked once, when system load exceeds the LoadWarningLevel for the LoadTimeThreshold. It is not invoked again until the LoadTimeCounter, which determines how many seconds system load has been above LoadWarningLevel, is reset.

About system limits and service group prerequisites

Limits is a system attribute and designates which resources are available on a system, including shared memory segments and semaphores.

Prerequisites is a service group attribute and helps manage application requirements. For example, a database may require three shared memory segments and 10 semaphores. VCS Load policy determines which systems meet the application criteria and then selects the least-loaded system.

If the prerequisites defined for a service group are not met on a system, the service group cannot be brought online on the system.

When configuring these attributes, define the service group's prerequisites first, then the corresponding system limits. Each system can have a different limit and there is no cap on the number of group prerequisites and system limits. Service group prerequisites and system limits can appear in any order.

You can also use these attributes to configure the cluster as N-to-1 or N-to-N. For example, to ensure that only one service group can be online on a system at a time, add the following entries to the definition of each group and system:

```
Prerequisites = { GroupWeight = 1 }
Limits = { GroupWeight = 1 }
```

System limits and group prerequisites work independently of FailOverPolicy. Prerequisites determine the eligible systems on which a service group can be started. When a list of systems is created, HAD then follows the configured FailOverPolicy.

Using capacity and limits

When selecting a node as a failover target, VCS selects the system that meets the service group's prerequisites and has the highest available capacity. If multiple systems meet the prerequisites and have the same available capacity, VCS selects the system appearing lexically first in the SystemList.

Systems having an available capacity of less than the percentage set by the LoadWarningLevel attribute, and those remaining at that load for longer than the time specified by the LoadTimeThreshold attribute invoke the loadwarning trigger.

Sample configurations depicting workload management

This section lists some sample configurations that use the concepts described in this chapter.

System and Service group definitions

The main.cf in this example shows various Service Group Workload Management attributes in a system definition and a service group definition.

See "VCS attributes" on page 701.

```
include "types.cf"
cluster SGWM-demo (
)

system LargeServer1 (
    Capacity = 200
    Limits = { ShrMemSeg=20, Semaphores=10, Processors=12 }
    LoadWarningLevel = 90
    LoadTimeThreshold = 600
    )

group G1 (
    SystemList = { LargeServer1, LargeServer2, MedServer1,
            MedServer2 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
            MedServer1=1, MedServer2=1 }
    AutoStartPolicy = Load
    AutoStartList = { MedServer1, MedServer2 }
    FailOverPolicy = Load
    Load = 100
    Prerequisites = { ShrMemSeg=10, Semaphores=5, Processors=6 }
    )
```

Sample configuration: Basic four-node cluster

```
include "types.cf"
cluster SGWM-demo

system Server1 (
    Capacity = 100
    )

system Server2 (
    Capacity = 100
    )

system Server3 (
    Capacity = 100
    )

system Server4 (
    Capacity = 100
    )

group G1 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 20
    )

group G2 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 40
    )

group G3 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 30
    )
```

```
group G4 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 10
    )

group G5 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 50
    )

group G6 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 30
    )

group G7 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 20
    )

group G8 (
    SystemList = { Server1, Server2, Server3, Server4 }
    AutoStartPolicy = Load
    AutoStartList = { Server1, Server2, Server3, Server4 }
    FailOverPolicy = Load
    Load = 40
    )
```

AutoStart operation

In this configuration, assume that groups probe in the same order they are described, G1 through G8. Group G1 chooses the system with the highest AvailableCapacity value. All systems have the same available capacity, so G1 starts on Server1 because this server is lexically first. Groups G2 through G4 follow on Server2 through Server4. With the startup decisions made for the initial four groups, the cluster configuration resembles:

Server	AvailableCapacity	Online Groups
Server1	80	G1
Server2	60	G2
Server3	70	G3
Server4	90	G4

As the next groups come online, group G5 starts on Server4 because this server has the highest AvailableCapacity value. Group G6 then starts on Server1 with AvailableCapacity of 80. Group G7 comes online on Server3 with AvailableCapacity of 70 and G8 comes online on Server2 with AvailableCapacity of 60.

The cluster configuration now resembles:

Server	AvailableCapacity	Online Groups
Server1	50	G1 and G6
Server2	20	G2 and G8
Server3	50	G3 and G7
Server4	40	G4 and G5

In this configuration, Server2 fires the loadwarning trigger after 600 seconds because it is at the default LoadWarningLevel of 80 percent.

Failure scenario

In the first failure scenario, Server4 fails. Group G4 chooses Server1 because Server1 and Server3 have AvailableCapacity of 50 and Server1 is lexically first. Group G5 then comes online on Server3. Serializing the failover choice allows complete load-based control and adds less than one second to the total failover time.

Following the first failure, the configuration now resembles:

Server	AvailableCapacity	Online Groups
Server1	40	G1, G6, and G4
Server2	20	G2 and G8
Server3	0	G3, G7, and G5

In this configuration, Server3 fires the loadwarning trigger to notify that the server is overloaded. An administrator can then switch group G7 to Server1 to balance the load across groups G1 and G3. When Server4 is repaired, it rejoins the cluster with an AvailableCapacity value of 100, making it the most eligible target for a failover group.

Cascading failure scenario

If Server3 fails before Server4 can be repaired, group G3 chooses Server1, group G5 chooses Server2, and group G7 chooses Server1. This results in the following configuration:

Server	AvailableCapacity	Online Groups
Server1	-10	G1, G6, G4, G3, and G7
Server2	-30	G2, G8, and G5

Server1 fires the loadwarning trigger to notify that it is overloaded.

Sample configuration: Complex four-node cluster

The cluster in this example has two large enterprise servers (LargeServer1 and LargeServer2) and two medium-sized servers (MedServer1 and MedServer2). It has four service groups, G1 through G4, with various loads and prerequisites. Groups G1 and G2 are database applications with specific shared memory and semaphore requirements. Groups G3 and G4 are middle-tier applications with no specific memory or semaphore requirements.

```
include "types.cf"
cluster SGWM-demo (
)

system LargeServer1 (
Capacity = 200
Limits = { ShrMemSeg=20, Semaphores=10, Processors=12 }
LoadWarningLevel = 90
LoadTimeThreshold = 600
)

system LargeServer2 (
Capacity = 200
Limits = { ShrMemSeg=20, Semaphores=10, Processors=12 }
LoadWarningLevel=70
LoadTimeThreshold=300
)

system MedServer1 (
Capacity = 100
Limits = { ShrMemSeg=10, Semaphores=5, Processors=6 }
)

system MedServer2 (
Capacity = 100
Limits = { ShrMemSeg=10, Semaphores=5, Processors=6 }
)
```

```
group G1 (
SystemList = { LargeServer1, LargeServer2, MedServer1,
MedServer2 }
SystemZones = { LargeServer1=0, LargeServer2=0, MedServer1=1,
        MedServer2=1 }
AutoStartPolicy = Load
AutoStartList = { LargeServer1, LargeServer2 }
FailOverPolicy = Load
Load = 100
Prerequisites = { ShrMemSeg=10, Semaphores=5, Processors=6 }
)

group G2 (
SystemList = { LargeServer1, LargeServer2, MedServer1,
MedServer2 }
SystemZones = { LargeServer1=0, LargeServer2=0, MedServer1=1,
        MedServer2=1 }
AutoStartPolicy = Load
AutoStartList = { LargeServer1, LargeServer2 }
FailOverPolicy = Load
Load = 100
Prerequisites = { ShrMemSeg=10, Semaphores=5, Processors=6 }
)

group G3 (
SystemList = { LargeServer1, LargeServer2, MedServer1,
MedServer2 }
SystemZones = { LargeServer1=0, LargeServer2=0, MedServer1=1,
        MedServer2=1 }
AutoStartPolicy = Load
AutoStartList = { MedServer1, MedServer2 }
FailOverPolicy = Load
Load = 30
)

group G4 (
SystemList = { LargeServer1, LargeServer2, MedServer1,
MedServer2 }
SystemZones = { LargeServer1=0, LargeServer2=0, MedServer1=1,
        MedServer2=1 }
AutoStartPolicy = Load
AutoStartList = { MedServer1, MedServer2 }
FailOverPolicy = Load
Load = 20
)
```

AutoStart operation

In this configuration, the AutoStart sequence resembles:

G1—LargeServer1

G2—LargeServer2

G3—MedServer1

G4—MedServer2

All groups begin a probe sequence when the cluster starts. Groups G1 and G2 have an AutoStartList of LargeServer1 and LargeServer2. When these groups probe, they are queued to go online on one of these servers, based on highest AvailableCapacity value. If G1 probes first, it chooses LargeServer1 because LargeServer1 and LargeServer2 both have an AvailableCapacity of 200, but LargeServer1 is lexically first. Groups G3 and G4 use the same algorithm to determine their servers.

Normal operation

The configuration resembles:

Server	AvailableCapacity	CurrentLimits	Online Groups
LargeServer1	100	ShrMemSeg=10 Semaphores=5 Processors=6	G1
LargeServer2	100	ShrMemSeg=10 Semaphores=5 Processors=6	G2
MedServer1	70	ShrMemSeg=10 Semaphores=5 Processors=6	G3
MedServer2	80	ShrMemSeg=10 Semaphores=5 Processors=6	G4

Failure scenario

In this scenario, if LargeServer2 fails, VCS scans all available systems in group G2's SystemList that are in the same SystemZone and creates a subset of systems that meet the group's prerequisites. In this case, LargeServer1 meets all required Limits. Group G2 is brought online on LargeServer1. This results in the following configuration:

Server	AvailableCapacity	CurrentLimits	Online Groups
LargeServer1	0	ShrMemSeg=0 Semaphores=0 Processors=0	G1, G2
MedServer1	70	ShrMemSeg=10 Semaphores=5 Processors=6	G3
MedServer2	80	ShrMemSeg=10 Semaphores=5 Processors=6	G4

After 10 minutes (LoadTimeThreshold = 600) VCS fires the loadwarning trigger on LargeServer1 because the LoadWarningLevel exceeds 90 percent.

Cascading failure scenario

In this scenario, another system failure can be tolerated because each system has sufficient Limits to accommodate the service group running on its peer. If MedServer1 fails, its groups can fail over to MedServer2.

If LargeServer1 fails, the failover of the two groups running on it is serialized. The first group lexically, G1, chooses MedServer2 because the server meets the required Limits and has AvailableCapacity value. Group G2 chooses MedServer1 because it is the only remaining system that meets the required Limits.

Sample configuration: Server consolidation

The following configuration has a complex eight-node cluster running multiple applications and large databases. The database servers, LargeServer1, LargeServer2, and LargeServer3, are enterprise systems. The middle-tier servers running multiple applications are MedServer1, MedServer2, MedServer3, MedServer4, and MedServer5.

In this configuration, the database zone (system zone 0) can handle a maximum of two failures. Each server has Limits to support a maximum of three database service groups. The application zone has excess capacity built into each server.

The servers running the application groups specify Limits to support one database, even though the application groups do not run prerequisites. This allows a database to fail over across system zones and run on the least-loaded server in the application zone.

```
include "types.cf"
cluster SGWM-demo (
)

system LargeServer1 (
    Capacity = 200
    Limits = { ShrMemSeg=15, Semaphores=30, Processors=18 }
    LoadWarningLevel = 80
    LoadTimeThreshold = 900
    )

system LargeServer2 (
    Capacity = 200
    Limits = { ShrMemSeg=15, Semaphores=30, Processors=18 }
    LoadWarningLevel=80
    LoadTimeThreshold=900
    )

system LargeServer3 (
    Capacity = 200
    Limits = { ShrMemSeg=15, Semaphores=30, Processors=18 }
    LoadWarningLevel=80
    LoadTimeThreshold=900
    )

system MedServer1 (
    Capacity = 100
    Limits = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )
```

```
system MedServer2 (
    Capacity = 100
    Limits = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )

system MedServer3 (
    Capacity = 100
    Limits = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )

system MedServer4 (
    Capacity = 100
    Limits = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )
system MedServer5 (
    Capacity = 100
    Limits = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )

group Database1 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { LargeServer1, LargeServer2, LargeServer3 }
    FailOverPolicy = Load
    Load = 100
    Prerequisites = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )

group Database2 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { LargeServer1, LargeServer2, LargeServer3 }
    FailOverPolicy = Load
    Load = 100
    Prerequisites = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )
```

```
group Database3 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { LargeServer1, LargeServer2, LargeServer3 }
    FailOverPolicy = Load
    Load = 100
    Prerequisites = { ShrMemSeg=5, Semaphores=10, Processors=6 }
    )

group Application1 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { MedServer1, MedServer2, MedServer3,
MedServer4,
            MedServer5 }
    FailOverPolicy = Load
    Load = 50
    )

group Application2 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { MedServer1, MedServer2, MedServer3,
MedServer4,
            MedServer5 }
    FailOverPolicy = Load
    Load = 50
    )
```

```
group Application3 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { MedServer1, MedServer2, MedServer3,
MedServer4,
            MedServer5 }
    FailOverPolicy = Load
    Load = 50
    )

group Application4 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { MedServer1, MedServer2, MedServer3,
MedServer4,
            MedServer5 }
    FailOverPolicy = Load
    Load = 50
    )

group Application5 (
    SystemList = { LargeServer1, LargeServer2, LargeServer3,
            MedServer1, MedServer2, MedServer3, MedServer4,
MedServer5 }
    SystemZones = { LargeServer1=0, LargeServer2=0,
LargeServer3=0,
            MedServer1=1, MedServer2=1, MedServer3=1,
MedServer4=1,
            MedServer5=1 }
    AutoStartPolicy = Load
    AutoStartList = { MedServer1, MedServer2, MedServer3,
MedServer4,
            MedServer5 }
    FailOverPolicy = Load
    Load = 50
    )
```

AutoStart operation

Based on the preceding main.cf example, the AutoStart sequence resembles:

Database1	LargeServer1
Database2	LargeServer2
Database3	LargeServer3
Application1	MedServer1
Application2	MedServer2
Application3	MedServer3
Application4	MedServer4
Application5	MedServer5

Normal operation

The configuration resembles:

Server	AvailableCapacity	CurrentLimits	Online Groups
LargeServer1	100	ShrMemSeg=10 Semaphores=20 Processors=12	Database1
LargeServer2	100	ShrMemSeg=10 Semaphores=20 Processors=12	Database2
LargeServer3	100	ShrMemSeg=10 Semaphores=20 Processors=12	Database3
MedServer1	50	ShrMemSeg=5 Semaphores=10 Processors=6	Application1
MedServer2	50	ShrMemSeg=5 Semaphores=10 Processors=6	Application2
MedServer3	50	ShrMemSeg=5 Semaphores=10 Processors=6	Application3

Server	AvailableCapacity	CurrentLimits	Online Groups
MedServer4	50	ShrMemSeg=5 Semaphores=10 Processors=6	Application4
MedServer5	50	ShrMemSeg=5 Semaphores=10 Processors=6	Application5

Failure scenario

In the following example, LargeServer3 fails. VCS scans all available systems in the SystemList for the Database3 group for systems in the same SystemZone and identifies systems that meet the group's prerequisites. In this case, LargeServer1 and LargeServer2 meet the required Limits. Database3 is brought online on LargeServer1. This results in the following configuration:

Server	AvailableCapacity	CurrentLimits	Online Groups
LargeServer1	0	ShrMemSeg=5 Semaphores=10 Processors=6	Database1 Database3
LargeServer2	100	ShrMemSeg=10 Semaphores=20 Processors=12	Database2

In this scenario, further failure of either system can be tolerated because each has sufficient Limits available to accommodate the additional service group.

Cascading failure scenario

If the performance of a database is unacceptable with two database groups running on a single server, the SystemZones policy can help expedite performance. Failing over a database group into the application zone has the effect of resetting the group's preferred zone. For example, in the above scenario Database3 was moved to LargeServer1. The administrator could reconfigure the application zone to move two application groups to a single system. The database application can then be switched to the empty application server (MedServer1–MedServer5), which would put Database3 in Zone1 (application zone). If a failure occurs in Database3, the group selects the least-loaded server in the application zone for failover.

Chapter 14

The role of service group dependencies

- About service group dependencies
- Service group dependency configurations
- Group Dependency FAQs
- Linking service groups
- VCS behavior with service group dependencies

About service group dependencies

Service groups can be dependent on each other. The dependent group is the *parent* and the other group is the *child*. For example a finance application (parent) may require that the database application (child) is online before it comes online. While service group dependencies offer more features to manage application service groups, they create more complex failover configurations.

A service group may function both as a parent and a child. Veritas Cluster Server supports five levels of service group dependencies.

Dependency links

The dependency relationship between a parent and a child is called a *link*. The link is characterized by the dependency category, the location of the service groups, and the rigidity of dependency.

- A dependency may be *online*, or *offline*.
- A dependency may be *local*, *global*, or *remote*.
- A dependency may be *soft*, *firm*, or *hard* with respect to the rigidity of the constraints between parent and child service group.

You can customize the behavior of service groups by choosing the right combination of the dependency category, location, and rigidity.

Dependency categories: online or offline dependencies

Dependency categories determine the relationship of the parent group with the state of the child group.

Online group dependency	The parent group must wait for the child group to be brought online before it can start. For example, to configure a database application and a database service as two separate groups, specify the database application as the parent, and the database service as the child.
Offline group dependency	The parent group can be started only if the child group is offline and vice versa. This behavior prevents conflicting applications from running on the same system. For example, configure a test application on one system as the parent and the production application on another system as the child.

Dependency location: local, global, or remote dependencies

The relative location of the parent and child service groups determines whether the dependency between them is a local, global, or remote.

Local dependency	The parent group depends on the child group being online or offline on the same system.
Global dependency	An instance of the parent group depends on one or more instances of the child group being online on any system.
Remote dependency	An instance of parent group depends on one or more instances of the child group being online on any system other than the system on which the parent is online.

Dependency rigidity: soft, firm, or hard dependencies

The type of dependency defines the rigidity of the link between parent and child groups. A soft dependency means minimum constraints, whereas a hard dependency means maximum constraints

Soft dependency	Specifies the minimum constraints while bringing parent and child groups online. The only constraint is that the child group must be online before the parent group is brought online. For example, in an online local soft dependency, an instance of the child group must be online on the same system before the parent group can come online. Soft dependency provides the following flexibility: ■ If the child group faults, VCS does not immediately take the parent offline. If the child group cannot fail over, the parent remains online. ■ When both groups are online, either group, child or parent, may be taken offline while the other remains online. ■ If the parent group faults, the child group may remain online. ■ When the link is created, the child group need not be online if the parent is online. However, when both groups are online, their online state must not conflict with the type of link.
Firm dependency	Imposes more constraints when VCS brings the parent or child groups online or takes them offline. In addition to the constraint that the child group must be online before the parent group is brought online, the constraints include: ■ If the child group faults, the parent is taken offline. If the parent is frozen at the time of the fault, the parent remains in its original state. If the child cannot fail over to another system, the parent remains offline. ■ If the parent group faults, the child group may remain online. ■ The child group cannot be taken offline if the parent group is online. The parent group can be taken offline while the child is online. ■ When the link is created, the parent group must be offline. However, if both groups are online, their online state must not conflict with the type of link.

Hard dependency	Imposes the maximum constraints when VCS brings the parent of child service groups online or takes them offline. For example: ■ If a child group faults, the parent is taken offline before the child group is taken offline. If the child group fails over, the parent fails over to another system (or the same system for a local dependency). If the child group cannot fail over, the parent group remains offline. ■ If the parent faults, the child is taken offline. If the child fails over, the parent fails over. If the child group cannot fail over, the parent group remains offline. **Note:** When the child faults, if the parent group is frozen, the parent remains online. The faulted child does not fail over. The following restrictions apply when configuring a hard dependency: ■ Only online local hard dependencies are supported. ■ Only a single-level, parent-child relationship can be configured as a hard dependency. ■ Only one parent and one child group can be configured in a hard dependency. ■ Bringing the child group online does not automatically bring the parent online. ■ Taking the parent group offline does not automatically take the child offline. ■ Bringing the parent online is prohibited if the child is offline.

Dependency limitations

- Multiple parent service groups may depend on a child service group, although a parent group may depend on only one child group.
- A group dependency tree may be at most five levels deep.

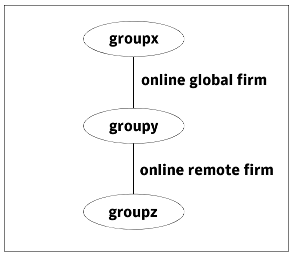

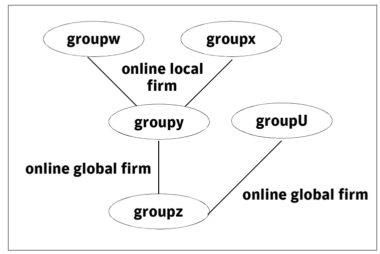

- You cannot link two service groups whose current states violate the relationship.

  For example, all link requests are accepted if all instances of parent group are offline.

  All link requests are rejected if parent group is online and child group is offline, except in offline dependencies.

  All online global/online remote link requests to link two parallel groups are rejected.

  All online local link requests to link a parallel parent group to a failover child group are rejected.

Service group dependency configurations

In the following tables, the term instance applies to parallel groups only. If a parallel group is online on three systems, for example, an instance of the group is online on each system. For failover groups, only one instance of a group is online at any time. The default dependency type is Firm.

Failover parent / Failover child

Link	Failover Parent Depends on …	Failover Parent is Online If …	If Failover Child Faults, then …	If Failover Parent Faults, then …
online local soft	Failover Child online on same system.	Child is online on same system.	Parent stays online. If Child fails over to another system, Parent migrates to the same system. If Child cannot fail over, Parent remains online.	Child stays online.
online local firm	Failover Child online on same system.	Child is online on same system.	Parent taken offline. If Child fails over to another system, Parent migrates to the same system. If Child cannot fail over, Parent remains offline.	Child stays online.
online local hard	Failover Child online on same system.	Child is online on same system.	Parents taken offline before Child is taken offline. If Child fails over to another system, Parent migrates to another system. If Child cannot fail 0ver, Parent remains offline.	Child taken offline. If Child fails over, Parent migrates to the same system. If Child cannot fail over, Parent remains offline.

Link	Failover Parent Depends on …	Failover Parent is Online If …	If Failover Child Faults, then …	If Failover Parent Faults, then …
online global soft	Failover Child online somewhere in the cluster.	Child is online somewhere in the cluster.	Parent stays online. If Child fails over to another system, Parent remains online. If Child cannot fail over, Parent remains online.	Child stays online. Parent fails over to any available system. If no system is available, Parent remains offline.
online global firm	Failover Child online somewhere in the cluster.	Child is online somewhere in the cluster.	Parent taken offline after Child is offlined. If Child fails over to another system, Parent is brought online on any system. If Child cannot fail over, Parent remains offline.	Child stays online. Parent fails over to any available system. If no system is available, Parent remains offline.
online remote soft	Failover Child online on another system.	Child is online on another system.	If Child fails over to the system on which Parent was online, Parent migrates to another system. If Child fails over to another system, Parent continues to run on original system. If Child cannot fail over, Parent remains online.	Child stays online. Parent fails over to a system where Child is not online. If the only system available is where Child is online, Parent is not brought online. If no system is available, Child remains online.

Link	Failover Parent Depends on …	Failover Parent is Online If …	If Failover Child Faults, then …	If Failover Parent Faults, then …
online remote firm	Failover Child online on another system.	Child is online on another system.	If Child fails over to the system on which Parent was online, Parent switches to another system. If Child fails over to another system, Parent restarts on original system. If Child cannot fail over, VCS takes the parent offline.	Parent fails over to a system where Child is not online. If the only system available is where Child is online, Parent is not brought online. If no system is available, Child remains online.
offline local	Failover Child offline on the same system	Child is offline on the same system.	If Child fails over to the system on which parent in not running, parent continues running. If child fails over to system on which parent is running, parent switches to another system, if available. If no system is available for Child to fail over to, Parent continues running.	Parent fails over to system on which Child is not online. If no system is available, Child remains online

Failover parent / Parallel child

With a failover parent and parallel child, no hard dependencies are supported.

Link	Failover Parent Depends on …	Failover Parent is Online if …	If Parallel Child Faults on a system, then …	If Failover Parent Faults, then …
online local soft	Instance of parallel Child group on same system.	Instance of Child is online on same system.	Parent fails over to other system and depends on Child instance there.	Parent fails over to other system and depends on Child instance there. Child Instance remains online where the Parent faulted.
online local firm	Instance of parallel Child group on same system.	Instance of Child is online on same system.	Parent is taken offline. Parent fails over to other system and depends on Child instance there.	Parent fails over to other system and depends on Child instance there. Child Instance remains online where Parent faulted.
online global soft	All instances of parallel Child group remaining online.	One or more instances of Child group is online somewhere in the cluster.	Parent remains online if Child faults on any system. If faulted Child fails over to another system, Parent is brought online on any system. If Child cannot fail over to another system, Parent remains offline.	Parent fails over to another system, maintaining dependence on all Child instances.

Link	Failover Parent Depends on …	Failover Parent is Online if …	If Parallel Child Faults on a system, then …	If Failover Parent Faults, then …
online global firm	All instances of parallel Child group remaining online.	All instances of Child group are online somewhere in the cluster.	Parent is taken offline. After Child fails over, Parent fails over to another system. If Child cannot fail over, Parent remains offline.	Parent fails over to another system, maintaining dependence on all Child instances.
online remote soft	One or more instances parallel Child group remaining online on other systems.	One or more instances of Child group are online on other systems.	Parent remains online. If Child fails over to the system on which Parent is online, Parent fails over to another system.	Parent fails over to another system, maintaining dependence on the Child instances.
online remote firm	All instances parallel Child group remaining online on other systems.	All instances of Child group are online on other systems.	Parent is taken offline. If Child fails over to the system on which Parent is online, Parent fails over to another system. If Child fails over to another system, Parent is brought online on its original system.	Parent fails over to another system, maintaining dependence on all Child instances.
offline local	Parallel Child offline on same system.	No instance of Child is online on same system.	Parent remains online if Child fails over to another system. If Child fails over to the system on which Parent is online, Parent fails over.	Child remains online.

Parallel parent / Failover child

Link	Parallel Parent Instances Depend on ...	Parallel Parent Instances are Online if ...	If Failover Child Faults on a system, then ...	If Parallel Parent Faults, then ...
online global soft	Failover Child group online somewhere in the cluster.	Failover Child is online somewhere in the cluster.	Parent remains online.	Child remains online
online global firm	Failover Child group somewhere in the cluster.	Failover Child is online somewhere in the cluster.	All instances of Parent taken offline. After Child fails over, Parent instances are brought failed over or restarted on the same systems.	Child stays online.
online remote soft	Failover Child group on another system.	Failover Child is online on another system.	If Child fails over to system on which Parent is online, Parent fails over to other systems. If Child fails over to another system, Parent remains online.	Child remains online. Parent tries to fail over to another system where child is not online.
online remote firm	Failover Child group on another system.	Failover Child is online on another system.	All instances of Parent taken offline. If Child fails over to system on which Parent was online, Parent fails over to other systems. If Child fails over to another system, Parent brought online on same systems.	Child remains online. Parent tries to fail over to another system where child is not online.

Link	Parallel Parent Instances Depend on …	Parallel Parent Instances are Online if …	If Failover Child Faults on a system, then …	If Parallel Parent Faults, then …
offline local	Failover Child offline on same system.	Failover Child is not online on same system.	Parent remains online if Child fails over to another system.	Child remains online.

Parallel parent / Parallel child

Global dependencies between parallel parent groups and parallel child groups are not supported.

Link	Parallel Parent Depends on …	Parallel Parent is Online If …	If Parallel Child Faults, then …	If Parallel Parent Faults, then …
online local soft	Parallel Child instance online on same system.	Parallel Child instance is online on same system.	If Child fails over to another system, Parent migrates to the same system as the Child. If Child cannot fail over, Parent remains online.	Child instance stays online. Parent instance can fail over only to system where Child instance is running and other instance of Parent is not running.
online local firm	Parallel Child instance online on same system.	Parallel Child instance is online on same system.	Parent taken offline. If Child fails over to another system, VCS brings an instance of the Parent online on the same system as Child. If Child cannot fail over, Parent remains offline.	Child stays online. Parent instance can fail over only to system where Child instance is running and other instance of Parent is not running.

Link	Parallel Parent Depends on …	Parallel Parent is Online If …	If Parallel Child Faults, then …	If Parallel Parent Faults, then …
offline local	Parallel Child offline on same system.	No instance of Child is online on same system.	Parent remains online if Child fails over to another system.	Child remains online.

Group Dependency FAQs

This section lists some commonly asked questions about group dependencies.

Dependency	Frequently asked questions
Online local	**Can child group be taken offline when parent group is online?** Soft=Yes Firm=No Hard = No. **Can parent group be switched while child group is online?** Soft=No Firm=No Hard = No. **Can child group be switched while parent group is online?** Soft=No Firm=No Hard = No.
Online global	**Can child group be taken offline when parent group is online?** Soft=Yes Firm=No. **Can parent group be switched while child group is running?** Soft=Yes Firm=Yes Hard=Yes. **Can child group be switched while parent group is running?** Soft=Yes Firm=No
Online remote	**Can child group be taken offline when parent group is online?** Firm=No Soft=Yes. **Can parent group be switched while child group is running?** Firm=Yes, but not to system on which child is running. Soft=Yes, but not to system on which child is running. **Can child group be switched while parent group is running?** Firm=No Soft=Yes, but not to system on which parent is running.
Offline local	**Can parent group be brought online when child group is offline?** Yes. **Can child group be taken offline when parent group is online?** Yes. **Can parent group be switched while the child group is running?** Yes, but not to system on which child is running. **Can child group be switched while the parent group is running?** Yes, but not to system on which parent is running.

Linking service groups

You can link service groups from the command line or from the Java and Web consoles.

Note that a configuration may require that a certain service group be running before another service group can be brought online. For example, a group containing resources of a database service must be running before the database application is brought online.

See also "Linking service groups" on page 193

To link service groups from the command line

◆ Type the following command

```
hagrp -link parent_group child_group gd_category
            gd_location gd_type
```

parent_group	Name of the parent group
child_group	Name of the child group
gd_category	category of group dependency (online/offline).
gd_location	the scope of dependency (local/global/remote).
gd_type	type of group dependency (soft/firm/hard). Default is firm.

VCS behavior with service group dependencies

VCS enables or restricts service group operations to honor service group dependencies. VCS rejects operations if the operation violates a group dependency.

Online operations in group dependencies

Typically, bringing a child group online manually is never rejected, except under the following circumstances:

- For online local dependencies, if parent is online, a child group online is rejected for any system other than the system where parent is online.
- For online remote dependencies, if parent is online, a child group online is rejected for the system where parent is online.
- For offline local dependencies, if parent is online, a child group online is rejected for the system where parent is online.

The following examples describe situations where bringing a parallel child group online is accepted:

- For a parallel child group linked online local with failover/parallel parent, multiple instances of child group online are acceptable.
- For a parallel child group linked online remote with failover parent, multiple instances of child group online are acceptable, as long as child group does not go online on the system where parent is online.
- For a parallel child group linked offline local with failover/parallel parent, multiple instances of child group online are acceptable, as long as child group does not go online on the system where parent is online.

Offline operations in group dependencies

VCS rejects offline operations if the procedure violates existing group dependencies. Typically, firm dependencies are more restrictive to taking child group offline while parent group is online. Rules for manual offline include:

- Parent group offline is never rejected.
- For all soft dependencies, child group can go offline regardless of the state of parent group.
- For all firm dependencies, if parent group is online, child group offline is rejected.
- For the online local hard dependency, if parent group is online, child group offline is rejected.

Switch operations in group dependencies

Switching a service group implies manually taking a service group offline on one system, and manually bringing it back online on another system. VCS rejects manual switch if the group does not comply with the rules for offline or online operations.

Section IV

Administration-Beyond the basics

- Chapter 13, "Controlling VCS behavior" on page 447
- Chapter 14, "The role of service group dependencies" on page 491
- Chapter 15, "VCS event notification" on page 511
- Chapter 16, "VCS event triggers" on page 527

Chapter 15

VCS event notification

- About VCS event notification
- Components of VCS event notification
- VCS events and traps
- Monitoring aggregate events
- Configuring notification

About VCS event notification

VCS provides a method for notifying important events such as resource or system faults to administrators or designated recipients. VCS includes a notifier component, which consists of the notifier process and the hanotify utility.

VCS support SNMP consoles that can use an SNMP V2 MIB.

The notifier process performs the following tasks:

- Receives notifications from HAD
- Formats the notification
- Generates an SNMP (V2) trap or sends an email to the designated recipient, or does both.

If you have configured owners for resources, groups, or for the cluster, VCS also notifies owners of events that affect their resources. A resource owner is notified of resource-related events, a group owner of group-related events, and so on.

See “VCS Attributes” for descriptions of the attributes that define owners for cluster objects.

There four severity levels: SevereError, Error, Warning, and Information. SevereError indicates the highest severity level, Information the lowest. Note that these severity levels are case-sensitive.

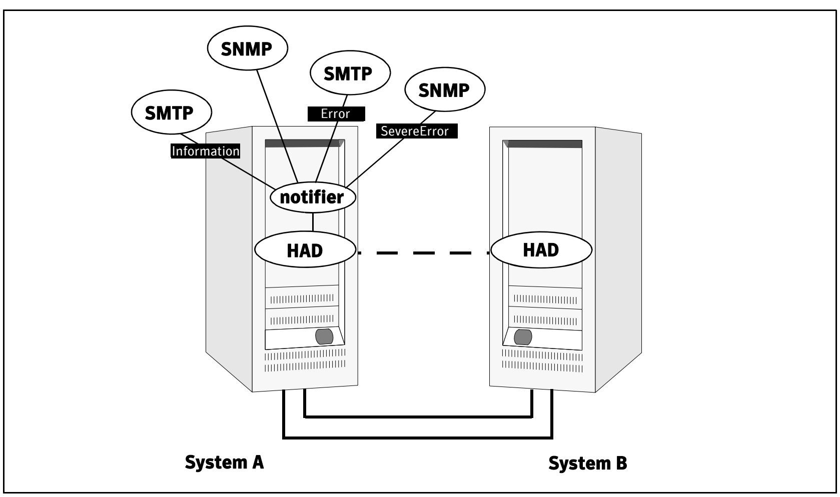

SNMP traps are forwarded to the SNMP console. Typically, traps are predefined for events such as service group or resource faults. You can use the hanotify utility to send additional traps.

Event messages and severity levels

When the VCS engine starts up, it queues all messages as Information. However, when notifier connects, it communicates one of the following severity levels to HAD, depending on which is the lowest:

- lowest severity for SNMP options
- lowest severity for SMTP options

If notifier is started from the command line without specifying a severity level for the SNMP console or SMTP recipients, notifier communicates the default severity level Warning to HAD. If notifier is configured under VCS control, severity must be specified. See the description of the NotifierMngr agent in the *Veritas Cluster Server Bundled Agents Reference Guide*.

For example, if the following severities are specified for notifier:

- Warning for email recipient 1
- Error for email recipient 2
- SevereError for SNMP console

Notifier communicates the minimum severity, Warning, to HAD, which then queues all messages labelled severity level Warning and greater.

Notifier ensures recipients gets only the messages they are designated to receive (according to the specified severity level). However, until notifier communicates the specifications to HAD, HAD stores all messages, because it does not know the severity the user has specified. This behavior prevents messages from being lost between the time HAD stores them and notifier communicates the specifications to HAD.

Persistent and replicated message queue

VCS includes a sophisticated mechanism for maintaining event messages, which ensures that messages are not lost. On each node, VCS queues messages to be sent to the notifier process. This queue is persistent as long as VCS is running and the contents of this queue remain the same on each node. If the notifier service group fails, notifier is failed over to another node in the cluster. Because the message queue is consistent across nodes, notifier can resume message delivery from where it left off even after failover.

How HAD deletes messages

The VCS engine, HAD, stores messages to be sent to notifier. HAD deletes messages under the following conditions:

- The message has been in the queue for one hour and notifier is unable to deliver the message to the recipient. (This behavior means that until notifier connects to HAD, messages are stored permanently in the queue until one of the following conditions are met.)
  or
- The message queue is full and to make room for the latest message, the earliest message is deleted.
  or
- VCS receives a message acknowledgement from notifier when notifier has delivered the message to at least one designated recipient.

  Example: two SNMP consoles and two email recipients are designated. Notifier sends an acknowledgement to VCS, even if the message reached only one of the four recipients. Error messages are also printed to the log files when delivery errors occur.

Components of VCS event notification

This section describes the notifier process and the hanotify utility.

The notifier process

The notifier process configures how messages are received from VCS and how they are delivered to SNMP consoles and SMTP servers. Using notifier, you can specify notification based on the severity level of the events generating the messages. You can also specify the size of the VCS message queue, which is 30 by default. You can change this value by modifying the MessageQueue attribute. See the *VCS Bundled Agents Reference Guide* for more information about this attribute.

When notifier is started from the command line, VCS does not control the notifier process. For best results, use the NotifierMngr agent that is bundled with VCS. Configure notifier as part of a highly available service group, which can then be monitored, brought online, and taken offline. For information about the agent, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

Note that notifier must be configured in a failover group, not parallel, because only one instance of notifier runs in the entire cluster. Also note that notifier does not respond to SNMP `get` or `set` requests; notifier is a trap generator only.

Notifier enables you to specify configurations for the SNMP manager and SMTP server, including machine names, ports, community IDs, and recipients' email addresses. You can specify more than one manager or server, and the severity level of messages that are sent to each.

Example of notifier command

```
notifier -s m=north -s m=south,p=2000,l=Error,c=your_company
      -t m=north,e="abc@your_company.com",l=SevereError
```

In this example, notifier:

- Sends all level SNMP traps to *north* at the default SNMP port and community value *public*.
- Sends Warning traps to *north*.
- Sends Error and SevereError traps to *south* at *port 2000* and community value *your_company*.
- Sends SevereError email messages to *north* as SMTP server at default port and to email recipient *abc@your_company.com*.

The hanotify utility

The hanotify utility enables you to construct user-defined messages. The utility forwards messages to HAD, which stores them in its internal message queue. Along with other messages, user-defined messages are also forwarded to the notifier process for delivery to email recipients, SNMP consoles, or both.

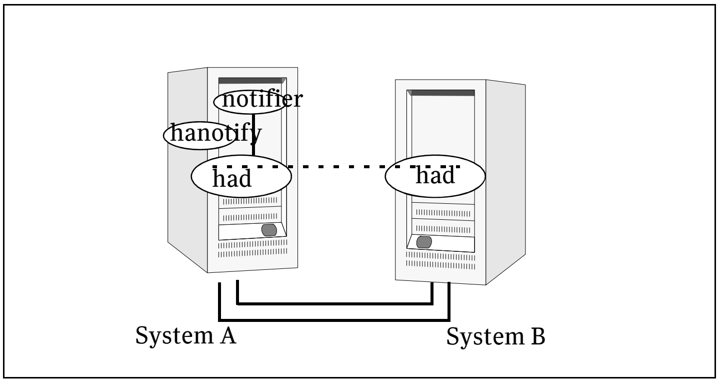

Example of hanotify command

```
hanotify -i 1.3.6.1.4.1.1302.3.8.10.2.8.0.10 -l Warning -n
      agentres -T 7 -t "custom agent" -o 4 -S sys1 -L mv -p
      sys2 -P mv -c MyAgent -C 7 -O johndoe -m "Custom message"
```

In this example, the number 1.3.6.1.4.1.1302.3.8.10.2.8.0.10 is the OID for the message being sent. Because it is a user-defined message, VCS has no way of knowing the OID associated with the SNMP trap corresponding to this message. Users must provide the OID.

The message severity level is set to Warning. The affected systems are sys1 and sys2. Running this command sends a custom message for the resource agentres from the agent MyAgent.

VCS events and traps

This section lists the events generate traps, email notification, or both. Note that SevereError indicates the highest severity level, Information the lowest. Traps specific to global clusters are ranked from Critical, the highest severity, to Normal, the lowest.

Events and traps for clusters

Event	Severity Level	Description
Cluster has faulted.	Error	Self-explanatory.
Heartbeat is down. (Global Cluster Option)	Error	The connector on the local cluster lost its heartbeat connection to the remote cluster.
Remote cluster is in RUNNING state. (Global Cluster Option)	Information	Local cluster has complete snapshot of the remote cluster, indicating the remote cluster is in the RUNNING state.
Heartbeat is "alive." (Global Cluster Option)	Information	Self-explanatory.
User has logged on to VCS.	Information	A user log on has been recognized because a user logged on by Cluster Manager, or because a `ha`*xxx* command was invoked.

Events and traps for agents

Event	Severity Level	Description
Agent is faulted.	Warning	The agent has faulted on one node in the cluster.
Agent is restarting	Information	VCS is restarting the agent.

Events and traps for resources

Event	Severity Level	Description
Resource state is unknown.	Warning	VCS cannot identify the state of the resource.
Resource monitoring has timed out.	Warning	Monitoring mechanism for the resource has timed out.
Resource is not going offline.	Warning	VCS cannot take the resource offline.
Health of cluster resource declined.	Warning	Used by agents to give additional information on the state of a resource. Health of the resource declined while it was online.
Resource went online by itself.	Warning (not for first probe)	The resource was brought online on its own.
Resource has faulted.	Error	Self-explanatory.
Resource is being restarted by agent.	Information	The agent is restarting the resource.
The health of cluster resource improved.	Information	Used by agents to give extra information about state of resource. Health of the resource improved while it was online.
Resource monitor time has changed.	Warning	This trap is generated when statistical analysis for the time taken by the monitor entry point of an agent is enabled for the agent. See "VCS agent statistics" on page 633. This trap is generated when the agent framework detects a sudden change in the time taken to run the monitor entry point for a resource. The trap information contains details of: ■ The change in time required to run the monitor entry point ■ The actual times that were compared to deduce this change.

Event	Severity Level	Description
Resource is in ADMIN_WAIT state.	Error	The resource is in the admin_wait state. See "Controlling Clean behavior on resource faults" on page 454.

Events and traps for systems

Event	Severity Level	Description
VCS is being restarted by `hashadow`.	Warning	Self-explanatory.
VCS is in jeopardy.	Warning	One node running VCS is in jeopardy.
VCS is up on the first node in the cluster.	Information	Self-explanatory.
VCS has faulted.	SevereError	Self-explanatory.
A node running VCS has joined cluster.	Information	Self-explanatory.
VCS has exited manually.	Information	VCS has exited gracefully from one node on which it was previously running.
CPU usage exceeded threshold on the system.	Warning	The system's CPU usage continuously exceeded the value that is set in the Notify threshold for a duration greater than the Notify time limit. See "When a resource comes online" on page 626.

Events and traps for service groups

Event	Severity Level	Description
Service group has faulted.	Error	Self-explanatory.
Service group concurrency violation.	SevereError	A failover service group has become online on more than one node in the cluster.
Service group has faulted and cannot be failed over anywhere.	SevereError	Specified service group faulted on all nodes where group could be brought online. There are no nodes to which the group can fail over.
Service group is online	Information	Self-explanatory.
Service group is offline.	Information	Self-explanatory.
Service group is autodisabled.	Information	VCS has autodisabled the specified group because one node exited the cluster.
Service group is restarting.	Information	Self-explanatory.
Service group is being switched.	Information	VCS is taking the service group offline on one node and bringing it online on another.
Service group restarting in response to persistent resource going online.	Information	Self-explanatory.
The global service group is online/partial on multiple clusters. (Global Cluster Option)	SevereError	A concurrency violation occurred for the global service group.
Attributes for global service groups are mismatched. (Global Cluster Option)	Error	The attributes ClusterList, AutoFailOver, and Parallel are mismatched for the same global service group on different clusters.

SNMP-specific files

VCS includes two SNMP-specific files: vcs.mib and vcs_trapd, which are created in /etc/VRTSvcs/snmp. The file vcs.mib is the textual MIB for built-in traps that are supported by VCS. Load this MIB into your SNMP console to add it to the list of recognized traps.

The file vcs_trapd is specific to the HP OpenView Network Node Manager (NNM) SNMP console. The file includes sample events configured for the built-in SNMP traps supported by VCS. To merge these events with those configured for SNMP traps:

```
xnmevents -merge vcs_trapd
```

When you merge events, the SNMP traps sent by VCS by way of notifier are displayed in the HP OpenView NNM SNMP console.

Note: For more information on xnmevents, see the HP OpenView documentation.

Trap variables in VCS MIB

Traps sent by VCS are reversible to SNMPv2 after an SNMPv2 -> SNMPv1 conversion.

For reversible translations between SNMPv1 and SNMPv2 trap PDUs, the second-last ID of the SNMP trap OID must be zero. This ensures that once you make a *forward* translation (SNMPv2 trap -> SNMPv1; RFC 2576 Section 3.2), the *reverse* translation (SNMPv1 trap --> SNMPv2 trap; RFC 2576 Section 3.1) is accurate.

The VCS notifier follows this guideline by using OIDs with second-last ID as zero, enabling reversible translations.

severityId

This variable indicates the severity of the trap being sent. It can take the following values:

Severity Level and Description	Value in Trap PDU
Information Important events exhibiting normal behavior	0
Warning Deviation from normal behavior	1
Error A fault	2
Severe Error Critical error that can lead to data loss or corruption	3

entityType and entitySubType

These variables specify additional information about the entity.

Entity Type	Entity Sub-type
Resource	String. For example, disk.
Group	The type of the group: ■ Failover ■ Parallel
System	String. For example, Solaris 2.8.
Heartbeat	The type of the heartbeat.
VCS	String
GCO	String
Agent name	Agent name

entityState

This variable describes the state of the entity.

Entity	States
VCS states	■ User has logged into VCS ■ Cluster has faulted ■ Cluster is in RUNNING state
Agent states	■ Agent is restarting ■ Agent has faulted
Resources states	■ Resource state is unknown ■ Resource monitoring has timed out ■ Resource is not going offline ■ Resource is being restarted by agent ■ Resource went online by itself ■ Resource has faulted ■ Resource is in admin wait state ■ Resource monitor time has changed

Entity	States
Service group states	■ Service group is online ■ Service group is offline ■ Service group is auto disabled ■ Service group has faulted ■ Service group has faulted and cannot be failed over anywhere ■ Service group is restarting ■ Service group is being switched ■ Service group concurrency violation ■ Service group is restarting in response to persistent resource going online ■ Service group attribute value does not match corresponding remote group attribute value ■ Global group concurrency violation
System states	■ VCS is up on the first node in the Cluster ■ VCS is being restarted by hashadow ■ VCS is in jeopardy ■ VCS has faulted ■ A node running VCS has joined cluster ■ VCS has exited manually ■ CPU Usage exceeded the threshold on the system
GCO heartbeat states	■ Cluster has lost heartbeat with remote cluster ■ Heartbeat with remote cluster is alive

Monitoring aggregate events

This section describes how you can detect aggregate events by monitoring individual notifications.

How to detect service group failover

VCS does not send any explicit traps when a failover occurs in response to a service group fault. When a service group faults, VCS generates the following notifications if the AutoFailOver attribute for the service group is set to 1:

- Service Group Fault for the node on which the service group was online and faulted
- Service Group Offline for the node on which the service group faulted
- Service Group Online for the node to which the service group failed over.

How to detect service group switch

When a service group is switched, VCS sends notification to indicate the following events:

- Service group is being switched
- Service Group Offline for the node from which the service group is switched
- Service Group Online for the node to which the service group was switched. This notification is sent after VCS completes the service group switch operation.

Note: You must configure appropriate severity for the notifier to receive these notifications. To receive VCS notifications, the minimum acceptable severity level is Information.

Detecting complementary events

Table 15-1 lists some events that complement each other, or cancel each other out.

Table 15-1 Complementary events in VCS

Event	Cancelling event
Remote cluster has faulted. (Global Cluster Option)	Remote cluster is in RUNNING state.
Heartbeat is down.	Heartbeat is alive.
Agent is faulted	Agent is restarting
Resource state is unknown.	Resource went online by itself.
Health of cluster resource declined.	Health of cluster resource improved.
VCS has faulted.	A node running VCS has joined cluster.
Service group has faulted.	Service group is online.
Service group is offline.	Service group is online
Service group is being switched.	Service group is online

Configuring notification

Configuring notification involves creating a resource for the Notifier Manager (NotifierMgr) agent in the ClusterService group. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the agent.

VCS provides several methods for configuring notification:

- Manually editing the main.cf file.
- Using the Notifier wizard.
  See "Setting up VCS event notification using the Notifier wizard" on page 227.

Chapter 16

VCS event triggers

- About VCS event triggers
- Using event triggers
- List of event triggers

About VCS event triggers

Triggers let you invoke user-defined scripts for specified events in a cluster.

VCS determines if the event is enabled and invokes the hatrigger script. The script is located at:

%VCS_HOME%\bin\hatrigger.pl

VCS also passes the name of the event trigger and associated parameters. For example, when a service group comes online on a system, VCS invokes the following command:

```
hatrigger -postonline system service_group.
```

VCS does not wait for the trigger to complete execution. VCS calls the trigger and continues normal operation.

VCS invokes event triggers on the system where the event occurred, with the following exceptions:

- VCS invokes the sysoffline and nofailover event triggers on the lowest-numbered system in the RUNNING state.
- VCS invokes the violation event trigger on all systems on which the service group was brought partially or fully online.

Using event triggers

VCS provides a sample Perl script for each event trigger at the following location:

%VCS_HOME%\bin\sample_triggers

Customize the scripts according to your requirements: you may choose to write your own Perl scripts.

To use an event trigger

1. Use the sample scripts to write your own custom actions for the trigger.
2. Move the modified trigger script to the following path on each node:
   %VCS_HOME%\bin\triggers
3. Configure other attributes that may be required to enable the trigger. See the usage information for the trigger for more information.
   See “List of event triggers” on page 529.

List of event triggers

The information in the following sections describes the various event triggers, including their usage, parameters, and location.

cpuusage event trigger

Description

The cpuusage event trigger is invoked when the system's CPU usage exceeds the ActionThreshold value of the system's CPUUsageMonitoring attribute for a duration longer than the ActionTimeLimit value. The trigger is not invoked if it was invoked on the system within the last five minutes.

See "When a resource comes online" on page 626.

This event trigger is configurable.

Usage

`- cpuusage` *triggertype system cpu_usage*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  If 0, the trigger is invoked from:
  %VCS_HOME%\bin\triggers\CPUUSage.extension (.exe, .pl,. ksh, or .bat)
  If 1, the system reboots by invoking the trigger from:
  %VCS_HOME%\bin\internal_triggers\cpuusage.pl.
- *system*—represents the name of the system.
- *cpu_usage*—represents the percentage of CPU utilization on the system.

To enable the trigger

Set following values in the system's CPUUsageMonitoring attribute:

- Enabled = 1
- ActionTimeLimit = Non-zero value representing time in seconds.
- ActionThreshold = Non-zero value representing CPU percentage utilization.
- Action = CUSTOM or REBOOT.
  CUSTOM—Invokes trigger from:
  %VCS_HOME%\bin\triggers\CPUUSage.extension (.exe, .pl,. ksh, or .bat)
  REBOOT—invokes trigger from:
  %VCS_HOME%\bin\internal_triggers\cpuusage.pl.
  and the system reboots.

To disable the trigger

Set one of the following values in CPUUsageMonitoring system attribute to 0 for the system:

- ActionTimeLimit = 0
- ActionThreshold = 0

injeopardy event trigger

Description Invoked when a system is in jeopardy. Specifically, this trigger is invoked when a system has only one remaining link to the cluster, and that link is a network link (LLT). This event is a considered critical because if the system loses the remaining network link, VCS does not fail over the service groups that were online on the system. Use this trigger to notify the administrator of the critical event. The administrator can then take appropriate action to ensure that the system has at least two links to the cluster.

This event trigger is non-configurable.

Usage `- injeopardy` *triggertype system system_state*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *system_state*—represents the value of the State attribute.

loadwarning event trigger

Description Invoked when a system becomes overloaded because the load of the system's online groups exceeds the system's LoadWarningLevel attribute for an interval exceeding the LoadTimeThreshold attribute.

For example, say the Capacity is 150, the LoadWarningLevel is 80, and the LoadTimeThreshold is 300. Also, the sum of the Load attribute for all online groups on the system is 135. Because the LoadWarningLevel is 80, safe load is 0.80*150=120. Actual system load is 135. If system load stays above 120 for more than 300 seconds, the LoadWarningLevel trigger is invoked.

Use this trigger to notify the administrator of the critical event. The administrator can then switch some service groups to another system, ensuring that no one system is overloaded.

This event trigger is non-configurable.

Usage `- loadwarning` *triggertype system available_capacity*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *available_capacity*—represents the system's AvailableCapacity attribute. (AvailableCapacity=Capacity-sum of Load for system's online groups.)

nofailover event trigger

Description Called from the lowest-numbered system in RUNNING state when a service group cannot fail over.

This event trigger is non-configurable.

Usage `- nofailover` *triggertype system service_group*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
- For this trigger, *triggertype*=0.
- *system*—represents the name of the last system on which an attempt was made to online the service group.
- *service_group*—represents the name of the service group.

postoffline event trigger

Description This event trigger is invoked on the system where the group went offline from a partial or fully online state. This trigger is invoked when the group faults, or is taken offline manually.

This event trigger is non-configurable.

Usage `-postoffline` *triggertype system service_group*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *service_group*—represents the name of the service group that went offline.

postonline event trigger

Description This event trigger is invoked on the system where the group went online from an offline state.

This event trigger is non-configurable.

Usage `-postonline` *triggertype system service_group*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *service_group*—represents the name of the service group that went online.

preonline event trigger

Description Indicates that when the HAD should call a user-defined script before bringing a service group online in response to the `hagrp -online` command or a fault.

If the trigger does not exist, VCS continues to bring the group online. If the script returns 0 without an exit code, VCS runs the `hagrp -online -nopre` command, with the `-checkpartial` option if appropriate.

If you do want to bring the group online, define the trigger to take no action. This event trigger is configurable.

Usage `-preonline` *triggertype system service_group whyonlining [system_where_group_faulted]*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *service_group*—represents the name of the service group on which the `hagrp` command was issued or the fault occurred.
- *whyonlining*—represents two values:
  FAULT: Indicates that the group was brought online in response to a group failover or switch.
  MANUAL: Indicates that group was brought online manually on the system that is represented by the variable *system*.
- *system_where_group_faulted*—represents the name of the system on which the group has faulted or switched. This variable is optional and set when the engine invokes the trigger during a failover or switch.

To enable the trigger Set the PreOnline attribute in the service group definition to 1.

To disable the trigger Set the PreOnline attribute in the service group definition to 0.

resadminwait event trigger

Description Invoked when a resource enters ADMIN_WAIT state.

When VCS sets a resource in the ADMIN_WAIT state, it invokes the resadminwait trigger according to the reason the resource entered the state.

See“Clearing resources in the ADMIN_WAIT state” on page 455.

This event trigger is non-configurable.

Usage `-resadminwait` *system resource adminwait_reason*

- *system*—represents the name of the system.
- *resource*—represents the name of the faulted resource.
- *adminwait_reason*—represents the reason the resource entered the ADMIN_WAIT state. Values range from 0-5:
  0 = The offline entry point did not complete within the expected time.
  1 = The offline entry point was ineffective.
  2 = The online entry point did not complete within the expected time.
  3 = The online entry point was ineffective.
  4 = The resource was taken offline unexpectedly.
  5 = The monitor entry point consistently failed to complete within the expected time.

resfault event trigger

Description Invoked on the system where a resource has faulted. Note that when a resource is faulted, resources within the upward path of the faulted resource are also brought down.

This event trigger is configurable.

To configure this trigger, you must define the following:

TriggerResFault: Set the attribute to 1 to invoke the trigger when a resource faults.

Usage `- resfault` *triggertype system resource previous_state*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *resource*—represents the name of the faulted resource.
- *previous_state*—represents the resource's previous state.

To enable the trigger To invoke the trigger when a resource faults, set the TriggerResFault attribute to 1.

resnotoff event trigger

Description Invoked on the system if a resource in a service group does not go offline even after issuing the offline command to the resource.

When invoked, the trigger script waits for a predefined interval and checks the state of the resource. If the resource is not offline, the trigger issues a system `shutdown` command, followed by the command `hastop -local -evacuate`.

This event trigger is configurable.

To configure this trigger, you must define the following:

Resource Name Define resources for which to invoke this trigger by entering their names in the following line in the script: `@resources = ("resource1","resource2");`

If any of these resources do not go offline, the trigger is invoked with that resource name and system name as arguments to the script.

$shutdown_timeout Define the time the script waits before it checks the resource state and issues a system shutdown command. For example, if this variable is set to 300, the script waits for 300 seconds before checking that the resource is offline and issuing the shutdown command.

$shutdown_countdown Define the time the script waits to shut down the system after issuing the `hastop -local -evacuate` command. For example, the value 300 indicates that the script waits for 300 seconds after issuing the `hastop -local -evacuate` command, and then shuts down the system.

Define this value to be greater than the time required to switch all service groups on the system to another system.

$forced_close_app Define whether the script forcefully closes all running applications when it triggers the system shutdown command. The value 1 indicates the script forcefully closes all running applications. The value 0 indicates it does not. Default is 1.

$reboot_option Define whether the script reboots the system after issuing the system shutdown command. The value 1 indicates the script reboots the system. The value 0 indicates it does not. Default is 1.

Usage `- resnotoff` *triggertype system resource*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the system on which the resource is not going offline.
- *resource*—represents the name of the resource.

resstatechange event trigger

Description

This trigger is invoked under the following conditions:

- Resource goes from OFFLINE to ONLINE.
- Resource goes from ONLINE to OFFLINE.
- Resource goes from ONLINE to FAULTED.
- Resource goes from FAULTED to OFFLINE. (When fault is cleared on non-persistent resource.)
- Resource goes from FAULTED to ONLINE. (When faulted persistent resource goes online or faulted non-persistent resource is brought online outside VCS control.)
- Resource is restarted by an agent because resource faulted and RestartLimit was greater than 0.

This event trigger is configurable.

Usage

`- resstatechange` *triggertype system resource previous_state new_state*

- *triggertype*—represents whether trigger is custom (*triggertype*=0) or internal (*triggertype*=1).
  For this trigger, *triggertype*=0.
- *system*—represents the name of the system.
- *resource*—represents the name of the resource.
- *previous_state*—represents the resource's previous state.
- *new_state*—represents the resource's new state.

To enable the trigger

This event trigger is not enabled by default. You must enable resstatechange by setting the attribute TriggerResStateChange to 1 in the main.cf file, or by issuing the command:

```
hagrp -modify service_group TriggerResStateChange 1
```

Note: Use the resstatechange trigger carefully. For example, enabling this trigger for a service group with 100 resources means 100 hatrigger processes and 100 resstatechange processes are fired each time the group is brought online or taken offline. Also, this is not a "wait-mode trigger. Specifically, VCS invokes the trigger and does not wait for trigger to return to continue operation

sysoffline event trigger

Description Called from the lowest-numbered system in RUNNING state when a system leaves the cluster.

This event trigger is non-configurable.

Usage `- sysoffline` *`system`* *`system_state`*

- *system*—represents the name of the system.
- *system_state*—represents the value of the State attribute.

See "System states" on page 696.

unable_to_restart_agent event trigger

Description This trigger is invoked when an agent faults more than a predetermined number of times with in an hour. When this occurs, VCS gives up trying to restart the agent. VCS invokes this trigger on the node where the agent faults.

You can use this trigger to notify the administrators that an agent has faulted, and that VCS is unable to restart the agent. The administrator can then take corrective action.

Usage `-unable_to_restart_agent` *`system`* *`resource_type`*

- *system*—represents the name of the system.
- *resource_type*—represents the resource type associated with the agent.

To disable the trigger Remove the files associated with the trigger from the *$VCS_HOME*/bin/triggers directory.

unable_to_restart_had event trigger

Description This event trigger is invoked by hashadow when hashadow cannot restart HAD on a system. If HAD fails to restart after six attempts, hashadow invokes the trigger on the system.

The default behavior of the trigger is to reboot the system. However, service groups previously running on the system are autodisabled when hashadow fails to restart HAD. Before these service groups can be brought online elsewhere in the cluster, you must autoenable them on the system. To do so, customize the unable_to_restart_had trigger to remotely execute the following command from any node in the cluster where VCS is running:

`hagrp -autoenable` *service_group* `-sys` *system*

For example, if hashadow fails to restart HAD on *system1*, and if *group1* and *group2* were online on that system, a trigger customized in this manner would autoenable *group1* and *group2* on *system1* before rebooting. Autoenabling *group1* and *group2* on *system1* enables these two service groups to come online on another system when the trigger reboots *system1*.

This event trigger is non-configurable.

Usage `-unable_to_restart_had`

This trigger has no arguments.

violation event trigger

Description This trigger is invoked only on the system that caused the concurrency violation. Specifically, it takes the service group offline on the system where the trigger was invoked. Note that this trigger applies to failover groups only. The default trigger takes the service group offline on the system that caused the concurrency violation.

This event trigger is non-configurable.

Usage `-violation` *system service_group*

- *system*—represents the name of the system.
- *service_group*—represents the name of the service group that was fully or partially online.

Section V

Multi-cluster configurations

- Chapter 17, "Connecting clusters–Creating global clusters" on page 543
- Chapter 17, "Administering global clusters from the Cluster Management Console" on page 537
- Chapter 19, "Administering global clusters from Cluster Manager (Java console)" on page 583
- Chapter 20, "Administering global clusters from the command line" on page 601
- Chapter 21, "Setting up replicated data clusters" on page 615

Chapter 17

Connecting clusters–Creating global clusters

- How VCS global clusters work
- VCS global clusters: The building blocks
- Prerequisites for global clusters
- Setting up a global cluster
- When a cluster faults
- Setting up a fire drill

How VCS global clusters work

Local clustering provides local failover for each site or building. But, these configurations do not provide protection against large-scale disasters such as major floods, hurricanes, and earthquakes that cause outages for an entire city or region. The entire cluster could be affected by an outage. In such situations, VCS global clusters ensure data availability can by migrating applications to remote clusters located considerable distances apart.

Let us take the example of an Oracle database configured in a VCS global cluster. Oracle is installed and configured in both clusters. Oracle data is located on shared disks within each cluster and is replicated across clusters to ensure data concurrency. The Oracle service group is online on a system in cluster A and is configured to fail over globally, on clusters A and B.

Figure 17-1 Sample global cluster setup

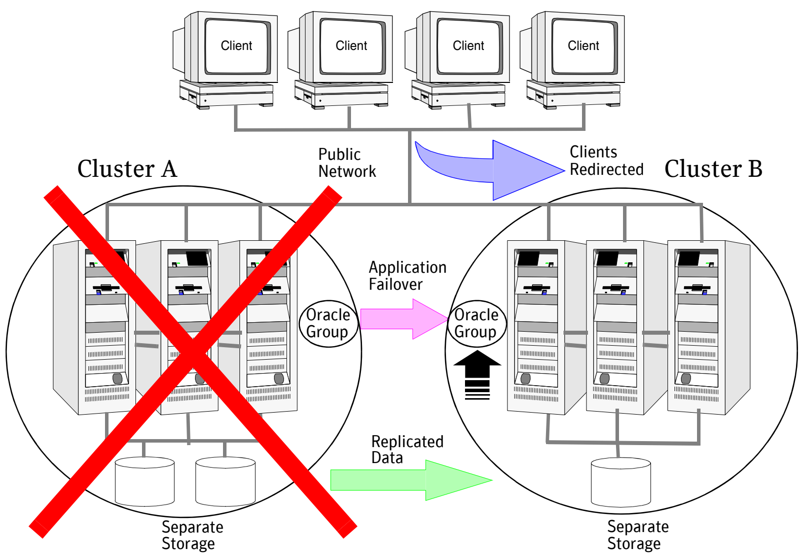

VCS continuously monitors and communicates events between clusters. Inter-cluster communication ensures that the global cluster is aware of the state of global service group at all times.

In the event of a system or application failure, VCS fails over the Oracle service group to another system in the same cluster. If the entire cluster fails, VCS fails over the service group to the remote cluster, which is part of the global cluster. VCS also redirects clients once the application is online on the new location.

VCS global clusters: The building blocks

VCS extends clustering concepts to wide-area high availability and disaster recovery with the following:

- Visualization of remote cluster objects
- Global service groups
- Global cluster management
- Serialization–The authority attribute
- Resiliency and “Right of way
- VCS framework
- The steward process: Split-brain in two-cluster global clusters

Visualization of remote cluster objects

VCS enables you to visualize remote cluster objects using the VCS command-line, the Java Console, and the Web Console.

You can define remote clusters in your configuration file, main.cf. The Remote Cluster Configuration wizard provides an easy interface to do so. The wizard updates the main.cf files of all connected clusters with the required configuration changes.

See “Adding a remote cluster” on page 585.

Global service groups

A *global* service group is a regular VCS group with additional properties to enable wide-area failover. The global service group attribute ClusterList defines the list of clusters to which the group can fail over. The service group must be configured on all participating clusters and must have the same name on each cluster. The Global Group Configuration wizard provides an easy interface to configure global groups.

See“Administering global service groups” on page 592.

Global cluster management

VCS enables you to perform operations (online, offline, switch) on global service groups from any system in any cluster using the VCS command-line interface, the Java Console, or the Web Console. You must log on with adequate privileges for cluster operations.

See "User privileges in global clusters" on page 66.

You can bring service groups online or switch them to any system in any cluster. If you do not specify a target system, VCS uses the FailOverPolicy to determine the system.

See "Defining failover policies" on page 452.

Management of remote cluster objects is aided by inter-cluster communication enabled by the wide-area connector (wac) process.

Wide-area connector process

The wide-area connector (wac) is a failover Process resource that ensures communication between clusters.

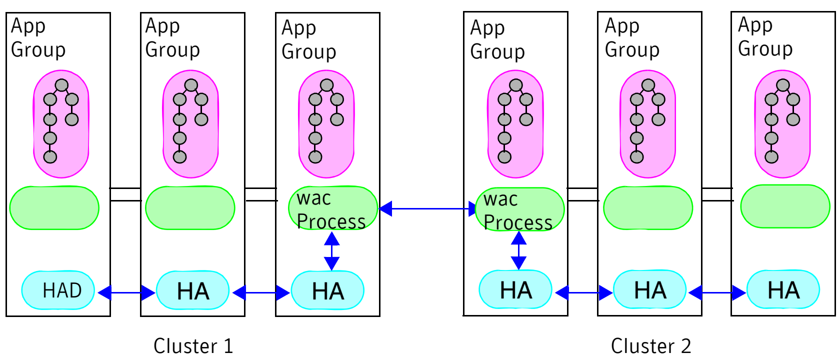

The wac process runs on one system in each cluster and connects with peers in remote clusters. It receives and transmits information about the status of the cluster, service groups, and systems. This communication enables VCS to create a consolidated view of the status of all the clusters configured as part of the global cluster. The process also manages wide-area heartbeating to determine the health of remote clusters. The process also transmits commands between clusters and returns the result to the originating cluster.

VCS provides the option of securing the communication between the wide-area connectors.

See "Secure communication in global clusters" on page 550.

Wide-area heartbeats

The wide-area Heartbeat agent manages the inter-cluster heartbeat. Heartbeats are used to monitor the health of remote clusters.

See "Heartbeat attributes" on page 742.

You can change the default values of the heartbeat agents using the command **`hahb -modify`**

Serialization–The authority attribute

VCS ensures that multi-cluster service group operations are conducted serially to avoid timing problems and to ensure smooth performance. The *Authority* attribute prevents a service group from coming online in multiple clusters at the same time. Authority is a persistent service group attribute and it designates which cluster has the right to bring a global service group online. The attribute cannot be modified at runtime.

A two-phase commit process prevents timing issues. If two administrators simultaneously try to bring a service group online in a two-cluster global group, one command is honored, and the other is rejected.

The attribute prevents bringing a service group online in a cluster that does not have the authority to do so. If the cluster holding authority is down, you can enforce a takeover by using the command `hagrp -online -force` *`service_group`*. This command enables you to fail over an application to another cluster when a disaster occurs.

Note: A cluster assuming authority for a group does not guarantee the group will be brought online on the cluster. The attribute merely specifies the right to attempt bringing the service group online in the cluster. The presence of Authority does not override group settings like frozen, autodisabled, non-probed, and so on, that prevent service groups from going online.

You must seed authority if it is not held on any cluster.

Offline operations on global groups can originate from any cluster and do not require a change of authority to do so, because taking a group offline does not necessarily indicate an intention to perform a cross-cluster failover.

Authority and AutoStart

The attributes Authority and AutoStart work together to avoid potential concurrency violations in multi-cluster configurations.

If the AutoStartList attribute is set, and if a group's Authority attribute is set to 1, HAD waits for the wac process to connect to the peer. If the connection fails, it means the peer is down and the AutoStart process proceeds. If the connection succeeds, HAD waits for the remote snapshot. If the peer is holding the authority for the group and the remote group is online (because of takeover), the local cluster does not bring the group online and relinquishes authority.

If the Authority attribute is set to 0, AutoStart is not invoked.

Resiliency and “Right of way

VCS global clusters maintain resiliency using the wide-area connector process and the ClusterService group. The wide-area connector process runs as long as there is at least one surviving node in a cluster.

The wide-area connector, its alias, and notifier are components of the ClusterService group, described in “The ClusterService Group” on page 13.

VCS framework

VCS agents now manage external objects that are part of wide-area failover. These objects include replication, DNS updates, and so on. These agents provide a robust framework for specifying attributes and restarts, and can be brought online upon fail over.

The steward process: Split-brain in two-cluster global clusters

Failure of all heartbeats between any two clusters in a global cluster indicates one of the following:

- The remote cluster is faulted.
- All communication links between the two clusters are broken.

In global clusters with more than three clusters, VCS queries the connected clusters to confirm that the remote cluster is truly down. This mechanism is called *inquiry*.

In a two-cluster setup, VCS uses the *Steward* process to minimize chances of a wide-area split-brain. The process runs as a standalone binary on a system outside of the global cluster configuration.

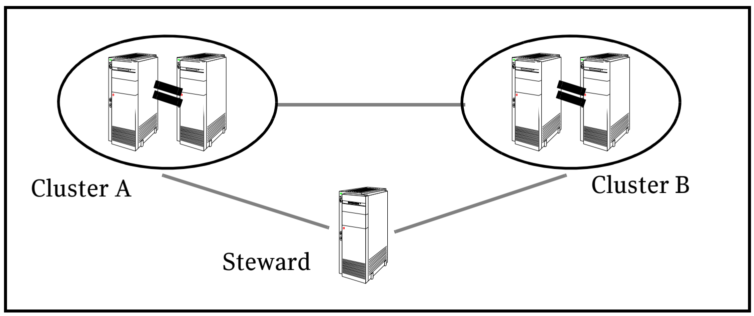

When all communication links between any two clusters are lost, each cluster contacts the Steward with an inquiry message. The Steward sends an ICMP ping to the cluster in question and responds with a negative inquiry if the cluster is running or with positive inquiry if the cluster is down. The Steward can also be used configurations with more than two clusters.

VCS provides the option of securing communication between the steward and the wide-area connectors.

See "Secure communication in global clusters" on page 550.

A Steward is effective only if there are independent paths from each cluster to the host running the Steward. If there is only one path between the two clusters, you must prevent split-brain by confirming manually via telephone or some messaging system with administrators at the remote site if a failure has occurred. By default, VCS global clusters fail over an application across cluster boundaries with administrator confirmation. You can configure automatic failover by setting the ClusterFailOverPolicy attribute to Auto.

See "Administering the cluster from the Cluster Management Console" on page 99.

Secure communication in global clusters

In global clusters, VCS provides the option of making the following communications secure:

- Communication between the wide-area connectors.
- Communication between the wide-area connectors and the Steward process.

For secure authentication, the wide-area connector process gets a security context as an account in the local authentication broker on each cluster node. The WAC account belongs to the same domain as HAD and Command Server and is specified as:

```
name = _WAC_GCO_(systemname)
domain = HA_SERVICES@(fully_qualified_system_name)
```

You must configure the wide-area connector process in all clusters to run in secure mode. If the wide-area connector process runs in secure mode, you must run the Steward in secure mode.

See “Configuring the Steward process (optional)” on page 559.

See “Prerequisites for secure clusters” on page 552.

Prerequisites for global clusters

This section describes the prerequisites for configuring global clusters.

Cluster setup

You must have at least two clusters to set up a global cluster. Every cluster must have the required licenses. A cluster can be part of one global cluster. VCS supports a maximum of four clusters participating in a global cluster.

Clusters must be running on the same platform; the operating system versions can be different. Clusters must be using the same VCS version.

Cluster names must be unique within each global cluster; system and resource names need not be unique across clusters. Service group names need not be unique across clusters; however, global service groups must have identical names.

Every cluster must have a valid virtual IP address, which is tied to the cluster. Define this IP address in the cluster's ClusterAddress attribute. This address is normally configured as part of the initial VCS installation. The IP address must have a DNS entry.

For remote cluster operations, you must configure a VCS user with the same name and privileges in each cluster.

See "User privileges in global clusters" on page 66.

Configured applications

Applications to be configured as global groups must be configured to represent each other in their respective clusters. The multiple application groups of a global group must have the same name in each cluster. The individual resources of the groups can be different. For example, one group might have a MultiNIC resource or more Mount-type resources. Clients redirected to the remote cluster in case of a wide-area failover must be presented with the same application they saw in the primary cluster.

However, the resources that make up a global group must represent the same application from the point of the client as its peer global group in the other cluster. Clients redirected to a remote cluster should not be aware that a cross-cluster failover occurred, except for some downtime while the administrator initiates or confirms the failover.

Wide-area heartbeats

There must be at least one wide-area heartbeat going from each cluster to every other cluster. VCS starts communicating with a cluster only after the heartbeat reports that the cluster is *alive*. VCS uses the ICMP ping by default, the infrastructure for which is bundled with the product. VCS configures the ICMP heartbeat if you use Cluster Manager (Java Console) to set up your global cluster. Other heartbeats must be configured manually.

ClusterService group

The ClusterService group must be configured with the Process (for the wide-area connector), NIC, and IP resources. The service group may contain additional resources for Cluster Management Console and notification, if these components are configured. It is configured automatically when VCS is installed or upgraded.

Replication setup

VCS global clusters are also used in case of disaster recovery, so you must set up real-time data replication between clusters. You can use VCS agents for supported replication solutions to manage the replication.

Prerequisites for secure clusters

- For both clusters to communicate in secure mode, both clusters must share a root broker or the root brokers must share a trust relationship.
- Both clusters must run in secure mode.
- If you plan to secure the communication between the wide-area connector processes, you must configure the processes in both clusters to run in secure mode.
- If the wide-area connector process runs in secure mode, you must run the Steward in secure mode.

Setting up a global cluster

This section describes the steps for planning, configuring, and testing a global cluster. It describes an example of converting a single instance Oracle database configured for local high availability in a VCS cluster to a highly available, disaster-protected infrastructure using a second cluster. The solution uses Veritas Volume Replicator to replicate changed data real-time.

In this example, an application is configured as a VCS service group (appgroup) on a two-node cluster.

Note: Before beginning the process, review the prerequisites listed in the section "Prerequisites for global clusters" on page 551 and make sure your configuration is ready for a global cluster application.

The process involves the following steps:

- Preparing the application for the global environment
- Configuring replication
- Linking the application and replication service groups
- Configuring the second cluster
- Linking clusters
- Configuring the Steward process (optional)
- Creating the global service group

Preparing the application for the global environment

Install the application in the second cluster. Make sure the installation is identical with the one in the first cluster.

Set up replication between the shared disk groups in both clusters. If your configuration uses VVR, the process involves grouping the shared data volumes in the first cluster into a Replicated Volume Group (RVG), and creating the VVR Secondary on hosts in the new cluster, located in your remote site.

Configuring the ClusterService group

You can configure the service group using the VCS Configuration wizard, Cluster Manager (Java Console), or the command line. For instructions on how to create the service group using the wizard, see "Configuring the ClusterService group" on page 394.

To configure the ClusterService group

1. If the ClusterService group does not exist in the cluster create a new service group called ClusterService.
2. Add resources of type IP, NIC, and Process to the service group.
3. Name the NIC resource csgnic and configure the following attribute for the resource:
   - MACAddress—The physical address of the adapter on the system. This attribute has a per-system value.
4. Name the IP resource webip and configure the following attributes for the resource:
   - MACAddress—The physical address of the adapter on the system. This attribute could have a per-system value.
   - Address—The virtual IP address for communicating between clusters. The IP address must have a DNS entry.
   - SubNetMask—The subnet mask associated with the virtual IP address.
5. Name the Process resource wac and configure the following attributes for the resource:
   - StartProgram—Complete path to the wide-area connector process.
     - If the clusters are running in secure mode, you can set this attribute to: %VCS_HOME%\bin\wac.exe -secure. For example: C:\Program Files\VERITAS\Cluster Server\bin\wac.exe -secure.
     - If the clusters are not running in secure mode, set this attribute to Secure: %VCS_HOME%\bin\wac.exe
       For example: C:\Program Files\VERITAS\Cluster Server\bin\wac.exe.
   - StopProgram—Complete path to the program that stops the wac process. Set this attribute to: %VCS_HOME%\bin\wacstop.exe
     For example: C:\Program Files\VERITAS\Cluster Server\bin\wacstop.exe.
   - MonitorProgram—Complete path to the program that monitors the wac process, typically C:\Program Files\VERITAS\Cluster Server\bin\wacmonitor.exe.
6. Mark the wac resource as critical.
7. Set resource dependencies as per the following information:
   - Process resource (wac) depends on the IP resource (webip)
   - IP resource (webip) depends on the NIC resource (csgnic)

   Enable the resources and bring the ClusterService group online.

Configuring replication

VCS supports several replication solutions for global clustering. Please contact your Symantec sales representative for the solutions supported by VCS. This section describes how to set up replication using Veritas Volume Replicator (VVR.)

Prerequisites

- Create Replicator Log Volumes for the primary and secondary sites.
- Create the replicated data sets for VVR. See the VVR documentation for instructions.
- Verify that the disk group is imported on the node on which you want to create the VVR RVG Service Group.
- Verify VCS is running, by running the following command on the host on which the you intend to run the VVR configuration Wizard.

To create a VVR service group

1. From the active node of the cluster at the primary site, click **Start>All Programs>Veritas>Veritas Cluster Server>Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
2. Read and verify the requirements on the Welcome panel, and click **Next.**
3. In the Wizard Options panel, click **Create a new replication service group**, and then click **Next**.
4. Specify the service group name and system priority list:
   - Enter the service group name.
   - In the Available Cluster Systems box, click the nodes on which to configure the service group, and click the right-arrow icon to move the nodes to the service group's system list. Make sure that the set of nodes selected for the replication service group is the same or a superset of nodes selected for the application's 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:
   - Select **Configure RVGPrimary resource for selected RVG**.
   - Select the replicated volume group for which you want to configure the VVR RVG resource.
   - Click **Next**.

7 In the IP Resource Options panel, select **Create a new IP resource** and click **Next**.

8 Enter the network information:
   - Verify or enter the virtual IP address; use the IP address specified as the primary IP address when you configured the RDS.
   - Verify the subnet mask.
   - Specify the adapters for each system in the configuration.
   - Click **Next**.

Note: At this step, the specified IP address does not yet need to exist.

9 If a message appears, indicating that the specified IP is not configured for replication in this RVG, click **OK** to continue.

10 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.
   - 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 in the right pane. Press Enter after editing each attribute. To cancel editing a resource name, press Esc.
   - Click **Next** to create the VVR service group.

11 When prompted, click **Yes** to create the service group.

Click **Finish** to bring the replication service group online.

Linking the application and replication service groups

Set an *online local hard* group dependency from appgroup to appgroup_rep to ensure that the service groups fail over and switch together.

To link the service groups

1. In the Cluster Explorer configuration tree, click the cluster name.
2. In the view panel, click the **Service Groups** tab. This opens the service group dependency graph.
3. Click **Link**.
4. Click the parent group, appgroup, and move the mouse toward the child group, appgroup_rep.
5. Click the child group appgroup_rep.
6. In the Link Service Groups dialog box, click the online local relationship and the firm dependency type and click **OK**.

Configuring the second cluster

1. Run the GCO Configuration wizard in the second cluster.
   See “Configuring the ClusterService group” on page 394.
2. Create a configuration that is similar to the one in the first cluster. You can do this by either using Cluster Manager (Java Console) to copy and paste resources from the primary cluster, or by copying the configuration of the appgroup and appgroup_rep groups from the main.cf file in the primary cluster to the secondary cluster.
   Run the VVR Configuration wizard to set up the VVR service group.
3. To assign remote administration privileges to users, configure users with the same name and privileges on both clusters.
   See “User privileges in global clusters” on page 66.
4. Make appropriate changes to the configuration. For example, you must modify the SystemList attribute to reflect the systems in the secondary cluster.
   Make sure that the name of the service group (appgroup) is identical in both clusters.
   It is a VVR best practice to use the same disk group and RVG name on both sites.
   If the volume names are the same on both sides, the Mount resources will mount the same block devices, and the same Oracle instance will start at the secondary in case of a failover.

Linking clusters

Once the VCS and VVR infrastructure has been set up at both sites, you must link the two clusters. The Remote Cluster Configuration wizard provides an easy interface to link clusters.

To link clusters

1. Verify the virtual IP address for the ClusterAddress attribute for each cluster is set. Use the same IP address as the one assigned to the IP resource in the ClusterService group.
2. If you are adding a stand-alone cluster to an existing global cluster environment, run the wizard from a cluster in the global cluster environment. Otherwise, run the wizard from any cluster. From Cluster Explorer, click Edit>Add/Delete Remote Cluster.
   See“Adding a remote cluster” on page 585.

To configure an additional heartbeat between the clusters (optional)

1. On Cluster Explorer’s **Edit** menu, click **Configure Heartbeats.**
2. In the Heartbeat configuration dialog box, enter the name of the heartbeat and select the check box next to the name of the cluster.
3. Click the icon in the **Configure** column to open the Heartbeat Settings dialog box.
4. Specify the value of the Arguments attribute and various timeout and interval fields. Click **+** to add an argument value; click **-** to delete it.
   If you specify IP addresses in the Arguments attribute, make sure the IP addresses have DNS entries.
5. Click **OK**.
6. Click **OK** in the Heartbeat configuration dialog box.
   Now, you can monitor the state of both clusters from the Java Console:

Configuring the Steward process (optional)

In case of a two-cluster GCO, you can configure a Steward to prevent potential split-brain conditions, provided the proper network infrastructure exists.

See "The steward process: Split-brain in two-cluster global clusters" on page 549.

To configure the Steward process for clusters not running in secure mode

1 Identify a system that will host the Steward process. Make sure both clusters can connect to the system through a ping command.

2 Copy the file steward from a node in the cluster to the Steward system. The file resides at the following path:
%VCS_HOME%\bin. The variable %VCS_HOME% represents the VCS installation directory, typically C:\Program Files\VERITAS\Cluster Server.

3 In both clusters, set the Stewards attribute to the IP address of the system running the Steward process. For example:

```
cluster cluster1938 (
UserNames = { admin = gNOgNInKOjOOmWOiNL }
ClusterAddress = "10.182.147.19"
Administrators = { admin }
CredRenewFrequency = 0
CounterInterval = 5
Stewards = {"10.212.100.165"}
}
```

4 On the system designated to host the Steward, start the Steward process:

```
steward.exe -start
```

To configure the Steward process for clusters running in secure mode

1 Verify the prerequisites for securing Steward communication are met.
See "Prerequisites for secure clusters" on page 552.

2 Identify a system that will host the Steward process. Make sure both clusters can connect to the system through a ping command.

3 Copy the steward file from a node in the cluster to the Steward system. The file resides at the following path:
%VCS_HOME%\bin\. The variable %VCS_HOME% represents the VCS installation directory, typically C:\Program Files\Veritas\Cluster Server.

4 Install the Symantec Product Authentication Services client on the system. See the *Quick Start Guide for Symantec Product Authentication Service* for instructions.

5 Create an account for the Steward in any authentication broker of the clusters that are part of the global cluster. All cluster nodes are authentication brokers when the cluster runs in secure mode.

```
vssat addprpl --pdrtype ab --domain
HA_SERVICES@<fully_qualified_name_of_cluster_node_on_which_t
his_command_is_being_run> --prplname Steward_GCO_systemname
--password password --prpltype service
```

When creating the account, make sure the following conditions are met:

- The domain name must be of the form: HA_SERVICES@*fully_qualified_system_name*
- The account name must be of the form: Steward_GCO_*systemname*
- The account type must be service and the domain type must be VX.

6 Note the password used to create the account.

7 Retrieve the broker hash for the account.

```
vssat showbrokerhash
```

8 Create a credential package (steward.cred) for this account. Note that the credential package will be bound to a host.

```
vssat createpkg --prplname Steward_GCO_systemname --domain
vx:HA_SERVICES@<fully_qualified_name_of_cluster_node_on_whic
h_this_command_is_being_run> --broker systemname:2821 --
password password --hash <brokerhash_obtained_in_above_step>
--out steward.cred --host_ctx
systemname_on_which_steward_will_run
```

9 Copy the file steward.cred to the C:\temp directory of the system designated to run the Steward process.

10 Execute the credential package on the system designated to run the Steward process.

```
vssat execpkg --in C:\temp\steward.cred --ob --host_ctx
```

11 On the Steward system, create a file called Steward.conf and populate it with the following information:

```
broker=system_name
accountname=accountname
domain=HA_SERVICES@FQDN_of_system_that_issued_the_certificate
```

12 In both clusters, set the Stewards attribute to the IP address of the system running the Steward process. For example:

```
cluster cluster1938 (
UserNames = { admin = gNOgNInKOjOOmWOiNL }
ClusterAddress = "10.182.147.19"
Administrators = { admin }
CredRenewFrequency = 0
CounterInterval = 5
Stewards = {"10.212.100.165"}
}
```

13 On the system designated to run the Steward, start the Steward process:

```
steward.exe -start -secure
```

To stop the Steward process

When you start the Steward, the process does not release the command window. Stop the Steward process, by typing control+C in the command window or open another command window and run the command to stop the Steward process.

- To stop the Steward process not running in secure mode, open a new command window and run the following command:

  ```
  steward.exe -stop
  ```

- To stop the Steward process running in secure mode, open a new command window and run the following command:

  ```
  steward.exe -stop -secure
  ```

Creating the global service group

Configure the Oracle service group, appgroup, as a global group by running the Global Group Configuration wizard.

To create the global service group

1. In the service group tree of Cluster Explorer, right-click the application service group (appgroup)
2. Select Configure and Global from the menu.
3. Enter the details of the service group to modify (appgroup).
4. From the **Available Clusters** box, click the clusters on which the group can come online. The local cluster is not listed as it is implicitly defined to be part of the ClusterList. Click the right arrow to move the cluster name to the **ClusterList** box.
5. Select the policy for cluster failover:
   - **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 faults.
   - **Connected** enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.
6. Click **Next**.
7. Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster.
8. Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
9. Enter the user name and the password for the remote cluster and click **OK**.
10. Click **Next**.
11. Click **Finish**.
12. Save the configuration.
    The appgroup service group is now a global group and can be failed over between clusters.
    For remote cluster operations, you must configure a VCS user with the same name and privileges in each cluster.
    See “User privileges in global clusters” on page 66.

When a cluster faults

In the global cluster setup, consider a case where the primary cluster suffers a failure. The Oracle service group cannot fail over in the local cluster and must fail over globally, to a node in another cluster.

In this situation, VCS sends an alert indicating that the cluster is down.

An administrator can bring the group online in the remote cluster.

The RVGPrimary agent ensures that VVR volumes are made writable and the DNS agent ensures that name services are resolved to the remote site. The application can be started at the remote site.

Declaring the type of failure

If a disaster disables all processing power in your primary data center, heartbeats from the failover site to the primary data center fail. VCS sends an alert signalling cluster failure. If you choose to take action on this failure, VCS prompts you to declare the type of failure.

You can choose one of the following options to declare the failure:

- *Disaster*, implying permanent loss of the primary data center
- *Outage,* implying the primary may return to its current form in some time
- *Disconnect*, implying a split-brain condition; both clusters are up, but the link between them is broken
- *Replica*, implying that data on the takeover target has been made consistent from a backup source and that the RVGPrimary can initiate a takeover when the service group is brought online. This option applies to VVR environments only.

You can select the groups to be failed over to the local cluster, in which case VCS brings the selected groups online on a node based on the group's FailOverPolicy attribute. It also marks the groups as being OFFLINE in the other cluster. If you do not select any service groups to fail over, VCS takes no action except implicitly marking the service groups as offline in the failed cluster.

Switching the service group back to the primary

You can switch the service group back to the primary after resolving the fault at the primary site. Before switching the application to the primary site, you must resynchronize any changed data from the active Secondary site since the failover. This can be done manually through VVR or by running a VCS action from the RVGPrimary resource.

To switch the service group when the primary site has failed and the secondary did a takeover

1 In the **Service Groups** tab of the configuration tree, right-click the resource.

2 Click **Actions**.

3 Specify the details of the action:

   - From the **Action** list, choose fbsync.
   - Click the system on which to execute the action.
   - Click **OK**.

   This begins a fast-failback of the replicated data set. You can monitor the value of the ResourceInfo attribute for the RVG resource to determine when the resynchronization has completed.

4 Once the resynchronization completes, switch the service group to the primary cluster.

   - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
   - Click **Switch To**, and click **Remote switch**.
   - In the Switch global group dialog box, click the cluster to switch the group. Click the specific system, or click **Any System**, and click **OK**.

Setting up a fire drill

The Disaster Recovery Fire Drill procedure tests the fault-readiness of a configuration by mimicking a failover from the primary site to the secondary site. This procedure is done without stopping the application at the primary site and disrupting user access, interrupting the flow of replicated data, or causing the secondary to need resynchronization.

The initial steps to create a fire drill service group on the secondary site that closely follows the configuration of the original application service group and contains a point-in-time copy of the production data in the Replicated Volume Group (RVG). Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to fail over 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.

You must conduct a fire drill only at the Secondary site; do not bring the fire drill service group online on the node hosting the original application.

You can set an offline local dependency between the fire drill service group and the application service group to make sure a fire drill does not block an application failover.

For detailed instructions on how to set up a fire drill in using the Solutions Configurations Center, see the following documents:

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

Chapter 18

Administering DR configurations

This chapter includes the following topics:

- About creating global service groups
- Administering global clusters
- Administering global service groups
- Administering global heartbeats
- Running a fire drill
- Viewing fire drill logs

About creating global service groups

The process of creating a global cluster environment involves the following tasks:

- Creating a common service group on specified local clusters and on clusters at a remote location
- Making sure the common service group is capable of being brought online on the local and remote clusters
- Designating the remote clusters as failover targets for the common service group (adding remote clusters)
- Converting the service group that is common to all the clusters to a global service group (creating the global service group)

After you set up the global cluster environment, you can configure global cluster heartbeats to monitor the health of the failover target clusters. Use the VCS Management Console to create global service groups, to add and delete remote clusters, and to create and monitor cluster heartbeats.

Administering global clusters

Creating a global cluster environment requires the following conditions:

- The ClusterService group for all clusters must be properly configured for VCS global cluster operations. The csgnic, gcoip, and wac resources must be able to come online.
- The service group intended to serve as the global group has the same unique name on all clusters that participate in the global environment.
- The clusters all have different, unique names.
- The clusters must use the same version of VCS.
- The clusters must use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.
- All clusters must either be configured in secure mode or not configured in secure mode. Mixing of modes is not supported.

For remote cluster operations, you must configure a VCS user with the same name and privileges in each cluster.

See “User privileges in global clusters” in the *Veritas Cluster Server Administrator’s Guide*.

Use the Cluster:Summary, Cluster:Resources, or Group:Summary views to administer a global cluster.

To navigate to the Cluster:Summary view

1. On the main tab bar, click **Manage**.
2. On the secondary tab bar, click **Summary**.
3. In the **Cluster:Summary** view, choose a task from the task pane or select another cluster-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the cluster.

To navigate to the Cluster:Resources view

1. On the main tab bar, click **Manage**.
2. On the secondary tab bar, click **Resources**.
3. In the **Cluster:Resources** view, choose a task from the task pane or select another cluster-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the cluster.

To navigate to the Group:Summary view

1. On the main tab bar, click **Manage**.
2. On the secondary tab bar, click **Groups**.
3. In the **Cluster:Groups** view, in the **Groups Listing** table, click a linked service group name.
4. In the **Group:Summary** view, choose a task from the task pane or select another cluster-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the cluster.

VCS global cluster terminology

All clusters and service groups configured into aVCS global cluster environment can be correctly referred to as global clusters and global service groups. However, in task descriptions, the following distinctions are made between a global cluster or service group and the corresponding remote cluster or service group:

- Global refers to the clusters and service groups that are:
  - Local to you (at your location)
  - Configured into a global cluster relationship with specific failover target clusters at another location.
- Remote refers to the clusters and service groups that are:
  - At another location
  - Configured in to a global cluster relationship as failover targets for the local clusters and service groups .

The terms are switched if the perspective is from the remote location.

Adding a remote cluster

Use this procedure to create global clusters by linking standalone clusters.

- If you are creating a global cluster environment for the first time with two standalone clusters, run the operation from either of the clusters.
- If you are adding a standalone cluster to an existing global cluster environment, run the operation from a cluster already in the global cluster environment.

The following information is required for this procedure:

- The IP address of the cluster, the IP address of a system in the cluster, or the name of a system in the cluster being added to the configuration.

- The user name and password of the administrator for the cluster being added to the configuration.

Veritas 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 to a global environment

1. In the **Cluster:Summary** view, in the **Configuration** task panel, click **Add/ Delete Remote Cluster**.
2. In the **Remote Cluster Configuration** wizard, read the introductory information and then click **Next**.
3. In the **Configuration Options** dialog box, click **Add Cluster** and then click **Next**.
4. Do one of the following:
   - For nonsecure clusters
     In the **Connection Details** dialog box, specify the following details for the connection to the remote cluster and then click **Next**:
     - A name or address
       Enter the IP address of the cluster, the IP address of a cluster system, or the name of a cluster system.
     - The port
       Verify the port number. The default is 14141.
     - An administrator user name and password.
       Enter an administrator-level user name and password that is valid on the remote cluster.
   - For secure clusters
     In the **Connection Details** dialog box, specify the following details for the connection to the remote cluster and then click **Next**:
     - A name or address
       Enter the IP address of the cluster, the IP address of a cluster system, or the name of a cluster system.
     - The port
       Verify the port number. The default is 14141.
     - Authentication credentials
       Choose to connect to the remote cluster with the credentials used for the current cluster connection or enter new credentials. You

must specify the user name, password, domain, and domain type. If you have connected to the remote cluster before using this wizard, you can use the credentials from the previous connection.

5 Click **Finish**.

The cluster icon changes to indicate that the cluster is a global cluster.

Deleting a remote cluster

The VCS Management Console enables you to delete a remote cluster. This operation involves the following tasks:

- Taking the wac resource in the ClusterService group offline in the cluster to be removed from the global environment. For example, to delete cluster C2 from a global environment containing C1 and C2, log in to C2 and take the wac resource offline.
- Removing the name of the specified cluster (C2) from the cluster lists of all global groups. The VCS Management Console updates the cluster lists for heartbeats. Log in to the local cluster (C1) to complete this task using the Global Groups Wizard.
- Removing the cluster (C2) from the local cluster (C1) using the Cluster:Summary view on the local cluster (C1).

Note: You cannot delete a remote cluster if the cluster is part of a cluster list for global service groups or global heartbeats, or if the cluster is in the RUNNING, BUILD, INQUIRY, EXITING, or TRANSITIONING states.

To take the wac resource offline

1. In the **Cluster:Resources** view, in the **Resources Listing** table, click the linked name of the **wac** resource.
2. In the **Resource:Summary** view, in the **Operations** task panel, click **Offline**.
3. In the **Offline Resource** dialog box, select the system on which you want to take the **wac** resource offline from the drop-down menu.
4. Click **OK**.

To remove a cluster from the cluster list for a global group

1. In the **Group:Summary** view, in the **Configuration** task panel, click **Configure Global Group**.
2. In the **Global Group Configuration** wizard, read the introductory information and then click **Next**.

3 In the **Cluster List Configuration** dialog box, under **Selected Clusters**, click the name of the cluster that you want to remove. Click the **<** (left-arrow) button.
   This action removes the cluster from the cluster list for the selected global service group and places it back under **Available Clusters**. To remove all clusters, click the **<<** (double-left-arrow) button.

4 Select the failover policy for the global service group and then click **Next**:
   - **Manual** prevents the service group from automatically failing over to another cluster.
   - **Auto** enables the service group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster faults.
   - **Connected** enables the service group to automatically fail over to another cluster if it is unable to fail over within the cluster.

5 In this step, you update the cluster list of remaining instances of the selected global service group. To perform the update, you must first verify or supply the authentication credentials for each remaining global cluster in the list. The VCS Management Console can then connect to those clusters and update the lists.
   In the **Remote Cluster Configuration** dialog box, verify the required information for the remaining remote clusters and then click **Next**.
   To change authentication information, click a cluster name under E**xisting Clusters** and then enter the authentication information in the fields to the right. The requisite information in this dialog box varies depending upon whether or not the cluster is configured in secure mode (uses an authentication broker).

6 Click **Next**.

7 Click **Finish**.
   You must repeat this procedure for all global service groups that are configured on the global cluster to be removed.

To remove a remote cluster from the local cluster

1 In the **Cluster:Summary** view, in the **Configuration** task panel, click **Add/Delete Remote Cluster**.

2 In the **Remote Cluster Configuration** wizard, select the cluster to delete and then click **Next**.
   The cluster that you want to delete cannot be in a running state.

3 In this step, you update the global cluster list of available global clusters. To perform the update, you must first verify or supply the authentication

credentials for each remaining global cluster. The VCS Management Console can then connect to those clusters and update the list.

In the **Remote Cluster Configuration** dialog box, verify the required information for the remaining remote clusters and then click **Next**.

To change authentication information, click a cluster name under E**xisting Clusters** and then enter the authentication information in the fields to the right. The requisite information in this dialog box varies depending upon whether or not the cluster is configured in secure mode (uses an authentication broker).

4. Click **Next**.
5. Click **Finish**.

Administering global service groups

After connecting clusters in a global cluster environment, use the Global Group Configuration wizard to convert a local service group that is common to the global clusters to a global group. This wizard also enables you to convert global groups into local groups.

Administering global groups requires the following conditions:

- A group that is intended to serve as the global group must have the same name across all applicable clusters.
- You must know the user name and password for the administrator for each cluster in the configuration.

Use the Cluster:Summary, Cluster:Resources, or Group:Summary views to administer a global service group.

To navigate to the Cluster:Summary view

1. On the main tab bar, click **Manage**.
2. On the secondary tab bar, click **Summary**.
3. In the **Cluster:Summary** view, choose a task from the task pane or select another cluster-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the cluster.

To navigate to the Cluster:Resources view

1. On the main tab bar, click **Manage**.
2. On the secondary tab bar, click **Summary** and then click **Resources**.

3 In the **Cluster:Resources** view, choose a task from the task pane or select another cluster-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the cluster.

To navigate to the Group:Summary view

1 On the main tab bar, click **Manage**.

2 On the secondary tab bar, click **Groups**.

3 In the **Cluster:Groups** view, in the **Groups Listing** table, click a linked service group name.

4 In the **Group:Summary** view, choose a task from the task pane or select another service-group-level view using the tabs on the secondary tab bar. Each view contains information and tasks for administering the service group.

Converting local service groups to global service groups

Use the Global Group Configuration wizard to configure a service group in a local cluster as a global service group.

To convert a service group on a local cluster to a global service group

1 In the **Cluster:Summary** view, in the **Groups Listing** table, click the linked name of the service group that you want to convert.
This service group should already have been commonly configured on at least one local and one remote cluster.

2 In the **Group:Summary** view, in the **Configuration** task panel, click **Configure Global Group**.

3 In the **Global Group Configuration** wizard, read the introductory information and click **Next**.

4 In the **Cluster List Configuration** dialog box, under **Available Clusters**, select the clusters on which the global service group can come online. To select a cluster, click the right-arrow button to move the cluster name under **Selected Clusters**.

5 Select the policy for service group failover and then click **Next**:

- **Manual** prevents a service group from automatically failing over to another cluster.
- **Auto** enables a service group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster faults.

- **Connected** enables a service group to automatically fail over to another cluster if it is unable to fail over within the cluster.

6 In this step, you update the cluster list of remaining instances of the selected global service group. To perform the update, you must first verify or supply the authentication credentials for each remaining global cluster in the list. The VCS Management Console can then connect to those clusters and update the lists.

   In the **Remote Cluster Configuration** dialog box, verify the required information for the remaining remote clusters and then click **Next**.

   To change authentication information, click a cluster name under E**xisting Clusters** and then enter the authentication information in the fields to the right. The requisite information in this dialog box varies depending upon whether or not the cluster is configured in secure mode (uses an authentication broker).

7 Click **No** if you want the operation to be completed only if the wizard can connect to all selected clusters.

8 Click **Next**.

9 Click **Finish**.

Bringing a remote service group online

Manually put a remote service group into a responsive, functioning state. You must have the role of cluster administrator or service group administrator to bring a remote service group online.

You can bring a remote service group online on a specific system, or you can bring the service group online "anywhere" in the remote cluster. If you select the Anywhere option, the service group is brought online on the first available system in the remote cluster.

To bring a remote service group online

1. In the **Cluster:Groups** view, in the **Groups Listing** table, click the linked name of the offline service group that you want to bring online.
2. In the **Group:Summary** view, in the **Remote Operations** task panel, click **Online**.
3. In the **Online Service Group** dialog box, specify the following details for the task and then click **OK**:
   - The target cluster

     Select the cluster on which you want to bring the remote service group online from the **Select the cluster you want to online this group on** drop-down menu. The cluster choices are populated using the cluster list for the remote service group.

- The target system
  Select the system on which you want to bring the remote service group online from the **Select the system you want to online this group on** drop-down menu.
  The system choices are populated using the system list for the remote service group. The **Anywhere** option causes this task to try each system in the list until the service group is successfully brought online.

Taking a remote service group offline

Manually put a remote service group into an unresponsive, nonfunctioning state. You must have the role of cluster administrator or service group administrator to take a service group offline.

You can take a remote service group offline on a specific system, or you can take the service group offline "anywhere". If you select the Anywhere option, the service group is taken offline on the first system in the remote cluster on which it is found.

To take a remote service group offline

1. In the **Cluster:Groups** view, in the **Groups Listing** table, click the linked name of the online service group that you want to take offline.
2. In the **Group:Summary** view, in the **Remote Operations** task panel, click **Offline**.
3. In the **Offline Service Group** dialog box, specify the following details for the task and then click **OK**:
   - The target cluster
     Select the cluster on which you want to take the remote service group offline from the **Select the cluster you want to offline this group on** drop-down menu. The cluster choices are populated using the cluster list for the remote service group.
   - The target system
     Select the system on which you want to take the remote service group offline from the **Select the system you want to offline this group on** drop-down menu.
     The system choices are populated using the system list for the remote service group. The **Anywhere** option causes this task to take the service group offline on the first system it is found to be running.

Switching a remote service group

Take a service group offline on one system in a cluster and bring it online on another system using a single task.

To switch a service group to another system

1. In the **Cluster:Groups** view, in the **Groups Listing** table, click the linked name of the remote service group that you want to switch.
2. In the **Group:Summary** view, in the **Remote Operations** task panel, click **Switch**.
3. In the **Switch Service Group** dialog box, specify the following details for the task and then click **OK**:
   - The target cluster
     Select the cluster to which you want to switch the remote service group from the **Select the cluster you want to switch this group to** drop-down menu. The cluster choices are populated using the cluster list for the remote service group.
   - The target system
     Select the system to which you want to switch the remote service group from the **Select the system you want to online this group on** drop-down menu.
     The system choices are populated using the system list for the remote service group. The **Anywhere** option causes this task to try every other system in the list until the service group is successfully brought online.

Administering global heartbeats

Use cluster heartbeats to monitor the health of clusters configured into a global cluster environment. A heartbeat is a list of clusters to which the VCS Management Console listens. The console listens for repetitive signals that confirm that communications with each cluster are valid and that each cluster is in an online state. Global clustering requires a minimum of one heartbeat between clusters; you can add additional heartbeats as a precautionary measure.

Use the Cluster:Heartbeats view to manage heartbeats. The Cluster Heartbeats tab is available only if global service groups are configured.

To navigate to the Cluster:Heartbeats view

1. On the main tab bar, click **Management**.
2. On the secondary tab bar, click **Cluster Heartbeats**.

Adding a global heartbeat

Add a heartbeat to monitor the health of a remote cluster that is configured in a global cluster environment.

To add a global heartbeat

1 In the **Cluster:Heartbeats** view, in the **Configuration** task panel, click **Add Heartbeat**.

2 In the **Add Heartbeat** dialog box, specify the following details for the heartbeat and then click **OK**.
   - The heartbeat name
     Enter a name for the heartbeat that is indicative of the generating it.
   - The global clusters to add to the cluster list for the heartbeat.
     In the **Cluster List for Heartbeat** table, check the box preceding the line item for each cluster that you want to participate in the heartbeat. Clear the check box to exclude clusters from the heartbeat list.
   - The Arguments attribute value, interval, and timeout options.
     Click the **...** (edit) button in the **Settings** column to specify the value for the **Arguments** attribute, the timeout options, and the interval options. Click **Save** to exit these options and save your selections.

Deleting a global heartbeat

Delete a heartbeat that is monitoring the health of a remote cluster that is configured in a global cluster environment.

Note: You cannot delete the heartbeat if it is the only remaining heartbeat between a global and remote cluster.

To delete a global heartbeat

1 In the **Cluster:Heartbeats** view, in the **Heartbeats Listing** table, click the **X** button in the line item corresponding to the heartbeat that you want to delete.

2 Click **OK** to confirm that you want to delete the heartbeat.

Modifying a global heartbeat

Change the clusters that participate in a heartbeat and the heartbeat parameters at any time.

1 In the **Cluster:Heartbeats** view, in the **Heartbeats Listing** table, click the **...** (edit) button in the line item corresponding to the heartbeat that you want to modify.

2 In the **Edit Heartbeat** dialog box, specify the following details for the heartbeat and then click **OK**:

- The heartbeat name
  Enter a name for the heartbeat that is indicative of the generating it.
- The global clusters to add to the cluster list for the heartbeat.
  In the **Cluster List for Heartbeat** table, check the box preceding the line item for each cluster that you want to participate in the heartbeat. Clear the check box to exclude clusters from the heartbeat list.
- The Arguments attribute value, interval, and timeout options.
  Click the **...** (edit) button in the **Settings** column to specify the value for the **Arguments** attribute, the timeout options, and the interval options. Click **Save** to exit these options and save your isolations.

Running a fire drill

The Run fire drill task enables you to do the following:

- Verify that replication for an application is working correctly
- Verify that a secondary disaster recovery (DR) application service group can be brought online successfully.

Three cluster objects play a role in a fire drill:

You can run a fire drill from the VCS Management Console. The VCS Management Console supports fire drills in multi-cluster mode only.

Before you run a fire drill, you must do the following:

- Configure the local (primary) and remote (secondary) global groups
- Set up the replication for the storage at the primary and secondary sites
- Configure the fire drill group using the FDSETUP command line wizard.

See the *Veritas Cluster Server Implementation Guide* for information on configuring fire drills.

To run a fire drill from the VCS Management Console

1 On the navigation bar, click **Home**.

2 On the secondary tab bar, click **Clusters**.

3 In the Home:Clusters view, in the Clusters Listing table, click the name of the primary global cluster.

4 On the secondary tab bar, click **Groups**.

5 In the Cluster:Groups view, in the Groups Listing table, click the name of the primary global group.

6 In the Group:Summary view, in the Remote Operations task panel, click **Run fire drill**.
You can view results of the fire drill in the Cluster:Groups view, the Group:Summary view, and in the Group:Fire Drill Logs view.
See "About service group listing tables" on page 137.

Viewing fire drill logs

Running a fire drill creates fire drill logs. If a service group is configured with a fire drill group, a tab labeled Fire Drill Logs appears on the secondary tab bar in the Group:Summary view.

See "Running a fire drill" on page 187.

To view fire drill logs

1 On the navigation bar, click **Home**.

2 On the secondary tab bar, click **Clusters**.

3 In the Home:Clusters view, in the Clusters Listing table, click the name of a VCS global cluster.
The global cluster must contain a global service group (primary group) that is configured with a fire drill group at a secondary location.

4 On the secondary tab bar, click **Groups**.

5 In the Cluster:Groups view, in the Groups Listing table, click the name of the primary global group.

6 In the Group:Summary view, on the secondary tab bar, click **Fire Drill Logs**.
This tab contains VCS log messages about the fire drill group on the remote (secondary) cluster and the resources that belong to it.

Chapter 19

Administering global clusters from Cluster Manager (Java console)

- About global clusters
- Adding a remote cluster
- Deleting a remote cluster
- Administering global service groups
- Administering global heartbeats

About global clusters

The process of creating a global cluster environment involves creating a common service group for specified clusters, making sure all the service groups are capable of being brought online in the specified clusters, connecting the standalone clusters, and converting the service group that is common to all the clusters to a global service group. Use the console to add and delete remote clusters, create global service groups, and manage cluster heartbeats.

Creating a global cluster environment requires the following conditions:

- All service groups are properly configured and able to come online.
- The service group that will serve as the global group has the same unique name across all applicable clusters.
- The clusters must use the same version of VCS.
- The clusters must use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.

Through the Java Console, you can simulate the process of generating and clearing global cluster faults in an OFFLINE state. Use VCS Simulator to complete these operations.

"Predicting VCS behavior using VCS Simulator" on page 405.

For remote cluster operations, you must configure a VCS user with the same name and privileges in each cluster.

See "User privileges in global clusters" on page 66.

Adding a remote cluster

Cluster Explorer provides a wizard to create global clusters by linking standalone clusters. Command Center only enables you to perform remote cluster operations on the local cluster.

- If you are creating a global cluster environment for the first time with two standalone clusters, run the wizard from either of the clusters.
- 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 to a global cluster environment 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 Clusters**.
2. Review the required information for the **Remote Cluster Configuration Wizard** and click **Next**.
3. In the Wizard Options dialog box:
   - Click **Add Cluster**.
   - Click **Next**.
4. Enter the details of the new cluster:

If the cluster is not running in secure mode:

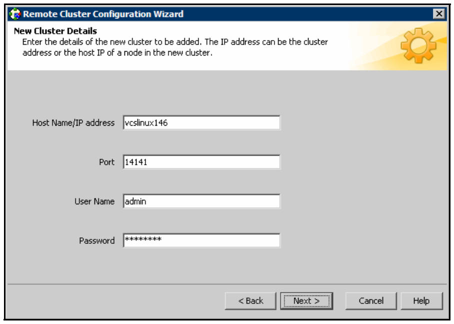

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Enter the user name and the password.
- Click **Next**.

If the cluster is running in secure mode:

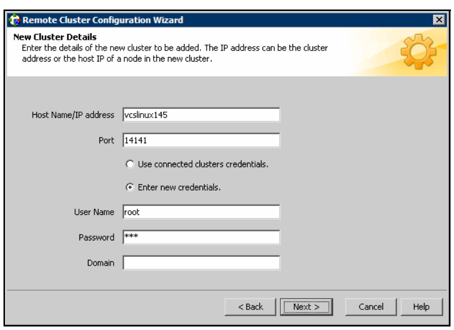

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Choose to connect to the remote cluster with the credentials used for the current cluster connection or enter new credentials, including the user name, password, and the domain.
  If you have connected to the remote cluster using the wizard earlier, you can use the credentials from the previous connection.

Click **Next**.

5 Enter the details of the existing remote clusters; this information on administrator rights enables the wizard to connect to all the clusters and make changes to the configuration:

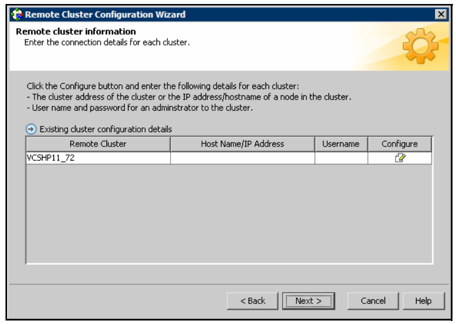

6 Click the **Configure** icon. The Remote cluster information dialog box is displayed.

If the cluster is not running in secure mode:

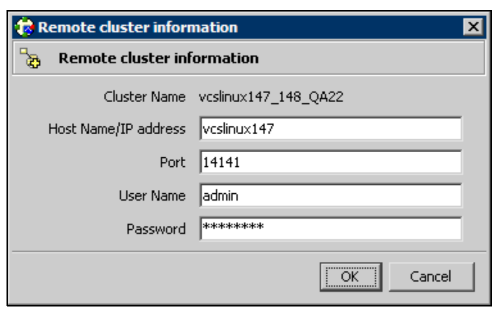

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Enter the user name.
- Enter the password.
- Click **OK**.
- Repeat these steps for each cluster in the global environment.

If the cluster is running in secure mode:

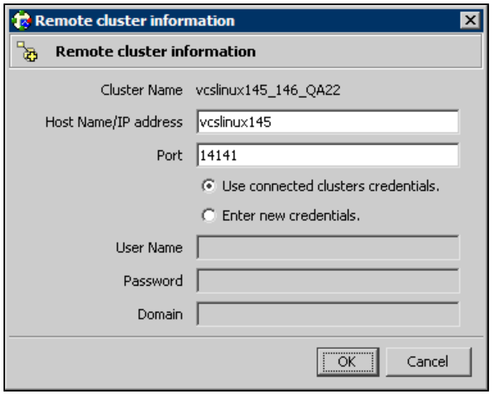

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Choose to connect to the remote cluster with the credentials used for the current cluster connection or enter new credentials, including the user name, password, and the domain.
- Click **OK**.

7 Click **Next**.

8 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are opened and changed; the wizard does not close the configurations.

To add a remote cluster to a global cluster environment in Command Center

---

Note: Command Center enables you to perform operations on the local cluster; this does not affect the overall global cluster configuration.

---

1 Click **Commands**>**Configuration**>**Cluster Objects**>**Add Remote Cluster**.

2 Enter the name of the cluster.

3 Enter the IP address of the cluster.

4 Click **Apply**.

Deleting a remote cluster

The Remote Cluster Configuration Wizard enables you to delete a remote cluster. This operation involves the following tasks:

- Taking the ApplicationProcess resource configured to monitor the wac resource offline on the cluster that will be removed from the global environment. For example, to delete cluster C2 from a global environment containing C1 and C2, log on to C2 and take the wac resource offline.
- Removing the name of the specified cluster (C2) from the cluster lists of the other global groups using the Global Group Configuration Wizard. Note that the Remote Cluster Configuration Wizard in Cluster Explorer updates the cluster lists for heartbeats. Log on to the local cluster (C1) to complete this task before using the Global Group Configuration Wizard.
- Deleting the cluster (C2) from the local cluster (C1) through the Remote Cluster Configuration Wizard.

Note: You cannot delete a remote cluster if the cluster is part of a cluster list for global service groups or global heartbeats, or if the cluster is in the RUNNING, BUILD, INQUIRY, EXITING, or TRANSITIONING states.

To take the wac resource offline

1. From Cluster Monitor, log on to the cluster that will be deleted from the global cluster environment.
2. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the **Process** type in the **ClusterService** group.
   or
   Click the ClusterService group in the configuration tree, click the **Resources** tab, and right-click the resource in the view panel.
3. Click **Offline**, and click the appropriate system from the menu.

To remove a cluster from a cluster list for a global group

1. From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
2. Click **Next**.
3. Enter the details of the service group to modify:
   - Click the name of the service group.
   - For global to local cluster conversion, click the left arrow to move the cluster name from the cluster list back to the **Available Clusters** box.
   - 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.

If the cluster is not running 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.
- Enter the user name.
- Enter the password.
- Click **OK**.

If the cluster is running 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 using the connected cluster's credentials or enter new credentials, including the user name, password, and the domain.
- Click **OK**.

5 Click **Next**.

6 Click **Finish**.

To delete a remote cluster from the local cluster

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 Clusters**.

2 Review the required information for the **Remote Cluster Configuration Wizard** and click **Next**.

3 In the Wizard Options dialog box:

- Click **Delete Cluster**.
- Click **Next**.

4 In the Delete Cluster dialog box:

- Click the name of the remote cluster to delete.
- Click **Next**.

5 Review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster.

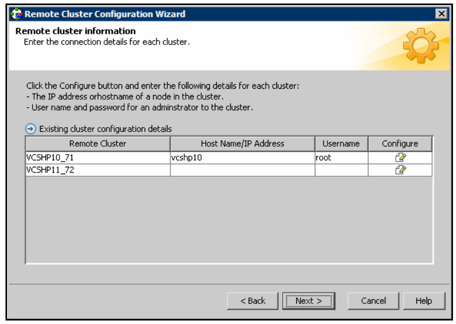

If the cluster is not running in secure mode:

- 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.
- Enter the user name.
- Enter the password.
- Click **OK**.

If the cluster is running 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 the domain.
  If you have connected to the remote cluster using the wizard earlier, you can use the credentials from the previous connection.
- Click **OK**.

6 Click **Finish**.

Administering global service groups

After connecting clusters in a global cluster environment, use the Global Group Configuration Wizard to convert a local service group that is common to the global clusters to a global group. This wizard also enables you to convert global groups into local groups.

Administering global groups requires the following conditions:

- A group that will serve as the global group must have the same name across all applicable clusters.
- You must know the user name and password for the administrator for each cluster in the configuration.

Use Cluster Explorer to bring a global group online and take a global group offline on a remote cluster.

Converting local and global groups

1. From Cluster Explorer, click **Configure Global Groups...** on the **Edit** menu.
   or
   From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global...** or **Make Local...** 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:

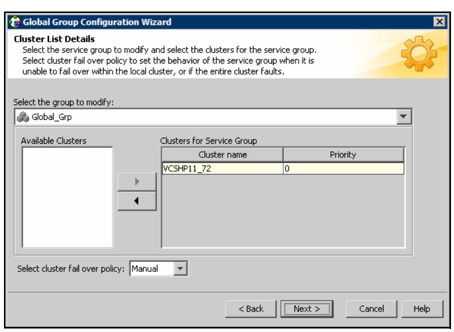

   - Click the name of the service group that will be converted from a local group to a global group, or vice versa.
   - From the **Available Clusters** box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the **Clusters for Service Group** box; 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 in which the group will attempt to come online. If necessary, double-click the entry in the **Priority** column to enter a new value.

- Select the policy for cluster failover:
  - **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 faults.
  - **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:

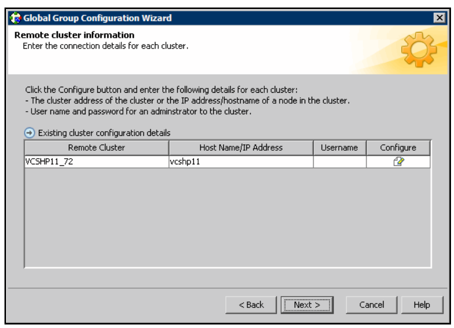

Click the **Configure** icon to review the remote cluster information for each cluster.

If the cluster is not running 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.
- Enter the user name and password.
- Click **OK**.

Repeat these steps for each cluster in the global environment.

If the cluster is running 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 the domain.
  If you have connected to the remote cluster using the wizard earlier, you can use the credentials from the previous connection.
- Click **OK**.

Repeat these steps for each cluster in the global environment.

5 In the Remote cluster information dialog box, click **Next**.

6 Click **Finish**.

Bringing a service group online in a remote cluster

1 In the **Service Groups** tab of the Cluster Explorer configuration tree of a local cluster, right-click the service group.
or
Click a local 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**.

4 In the Question dialog box, click **Yes**.

Taking a service group offline in a remote cluster

1. In the **Service Groups** tab of the Cluster Explorer configuration tree of a local cluster, right-click the service group.
   or
   Click a local cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
2. Click **Offline**, and click **Remote offline...**
3. In the Offline global group dialog box:
   - Click the remote cluster to take the group offline.
   - Click the specific system, or click **All Systems**, to take the group offline.
   - Click **OK**.
4. In the Question dialog box, click **Yes**.

Switching a service group to a remote cluster

1. In the **Service Groups** tab of the Cluster Explorer configuration tree of a local cluster, right-click the service group.
   or
   Click a local cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
2. Click **Switch To**, and click **Remote switch...**
3. In the Switch global group dialog box:
   - Click the cluster to switch the group.
   - Click the specific system, or click **Any System**, to take the group offline.
   - Click **OK**.
4. In the Question dialog box, click **Yes**.

Administering global heartbeats

Use Cluster Explorer to add, modify, and delete heartbeats in a global cluster environment. *Icmp* heartbeats send Icmp packets simultaneously to all IP addresses; *IcmpS* heartbeats send individual Icmp packets to IP addresses in serial order. Global clustering requires a minimum of one heartbeat between clusters; the Icmp heartbeat is added when the cluster is added to the environment. You can add additional heartbeats as a precautionary measure.

Adding a global heartbeat

To add a cluster heartbeat from Cluster Explorer

1 Click **Configure Heartbeats** on the **Edit** menu.

2 In the Heartbeat Configuration dialog box:

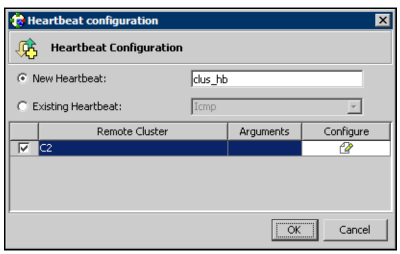

- Enter the name of the heartbeat.
- Select the check box next to the name of the cluster to add it to the cluster list for the heartbeat.
- Click the icon in the **Configure** column to open the Heartbeat Settings dialog box.

- Specify the value of the Arguments attribute and various timeout and interval fields. Click **+** to add an argument value; click **-** to delete it.

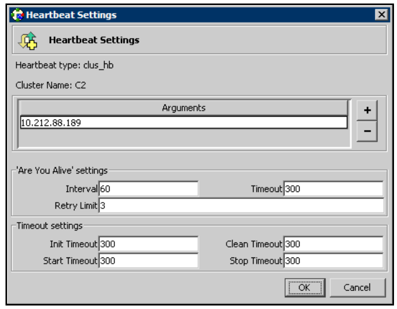

- Click **OK**.
- Click **OK** on the Heartbeat configuration dialog box.

To add a cluster heartbeat from Command Center

1. Click **Commands**>**Configuration**>**Cluster Objects**>**Add Heartbeat**.
2. Enter the name of the heartbeat.
3. Click **Apply**.

Modifying a global heartbeat

1 From Cluster Explorer, click **Configure Heartbeats** on the **Edit** menu.

2 In the Heartbeat Configuration dialog box:

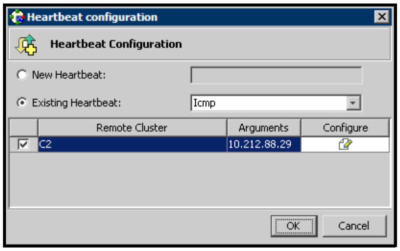

- Click **Existing Heartbeat**.
- Click the name of the existing heartbeat from the menu.
- Select or clear the check box next to the name of a cluster to add or remove it from the cluster list for the heartbeat.
- If necessary, click the icon in the **Configure** column to open the Heartbeat Settings dialog box. Otherwise, proceed to the last step.
- Change the values of the Arguments attribute and various timeout and interval fields. Click **+** to add an argument value; click **-** to delete it.

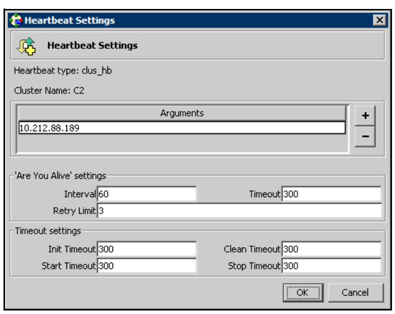

- Click **OK**.
- Click **OK** on the Heartbeat Configuration dialog box.

Deleting a global heartbeat

Note: You cannot delete the last heartbeat between global clusters.

To delete a cluster heartbeat from Command Center

1. Click **Commands**>**Configuration**>**Cluster Objects**>**Delete Heartbeat**.
2. Click the heartbeat to delete.
3. Click **Apply**.

Chapter 20

Administering global clusters from the command line

- About administering global clusters from the command line
- Global querying
- Administering global service groups
- Administering resources
- Administering clusters in global clusters
- Administering heartbeats

About administering global clusters from the command line

For remote cluster operations, you must configure a VCS user with the same name and privileges in each cluster.

See "User privileges in global clusters" on page 66.

Global querying

VCS enables you to query global cluster objects, including service groups, resources, systems, resource types, agents, and clusters. You may enter query commands from any system in the cluster. Commands to display information on the global cluster configuration or system states can be executed by all users; you do not need root privileges. Only global service groups may be queried.

Querying global cluster service groups

To display service group attribute values across clusters

```
hagrp -value service_group attribute [system] [-clus cluster |
        -localclus]
```

The option `-clus` displays the attribute value on the cluster designated by the variable *cluster*; the option `-localclus` specifies the local cluster.

If the attribute has local scope, you must specify the system name, except when querying the attribute on the system from which you run the command.

To display the state of a service group across clusters

```
hagrp -state [service_groups -sys systems] [-clus cluster |
        -localclus]
```

The option `-clus` displays the state of all service groups on a cluster designated by the variable *cluster*; the option `-localclus` specifies the local cluster.

To display service group information across clusters

```
hagrp -display [service_groups] [-attribute attributes]
     [-sys systems] [-clus cluster | -localclus]
```

The option `-clus` applies to global groups only. If the group is local, the cluster name must be the local cluster name, otherwise no information is displayed.

To display service groups in a cluster

```
hagrp -list [conditionals] [-clus cluster | -localclus]
```

The option `-clus` lists all service groups on the cluster designated by the variable *cluster*; the option `-localclus` specifies the local cluster.

To display usage for the service group command

```
hagrp [-help [-modify|-link|-list]]
```

Querying resources

To display resource attribute values across clusters

```
hares -value resource attribute [system] [-clus cluster |
       -localclus]
```

The option `-clus` displays the attribute value on the cluster designated by the variable *cluster*; the option `-localclus` specifies the local cluster.

If the attribute has local scope, you must specify the system name, except when querying the attribute on the system from which you run the command.

To display the state of a resource across clusters

```
hares -state [resource -sys system] [-clus cluster | -
localclus]
```

The option `-clus` displays the state of all resources on the specified cluster; the option `-localclus` specifies the local cluster. Specifying a system displays resource state on a particular system.

To display resource information across clusters

```
hares -display [resources] [-attribute attributes] [-group
       service_groups][-type types] [-sys systems] [-clus
cluster |
       -localclus]
```

The option `-clus` lists all service groups on the cluster designated by the variable *cluster*; the option `-localclus` specifies the local cluster.

For a list of resources across clusters

```
hares -list [conditionals] [-clus cluster | -localclus]
```

The option `-clus` lists all resources that meet the specified conditions in global service groups on a cluster as designated by the variable *cluster*.

To display usage for the resource command

```
hares -help [-modify | -list]
```

Querying systems

To display system attribute values across clusters

hasys -value *system attribute* [-clus *cluster* | -localclus]

The option `-clus` displays the values of a system attribute in the cluster as designated by the variable *cluster;* the option `-localclus` specifies the local cluster.

To display the state of a system across clusters

hasys -state [*system*] [-clus *cluster* | -localclus]

Displays the current state of the specified system. The option `-clus` displays the state in a cluster designated by the variable *cluster;* the option `-localclus` specifies the local cluster. If you do not specify a system, the command displays the states of all systems.

For information about each system across clusters

**hasys -display [*systems*] [-attribute *attributes*] [-clus *cluster* |
-localclus]**

The option `-clus` displays the attribute values on systems (if specified) in a cluster designated by the variable *cluster;* the option `-localclus` specifies the local cluster.

For a list of systems across clusters

hasys -list [*conditionals*] [-clus *cluster* | -localclus]

Displays a list of systems whose values match the given conditional statements. The option `-clus` displays the systems in a cluster designated by the variable *cluster;* the option `-localclus` specifies the local cluster.

Querying clusters

For the value of a specific cluster attribute on a specific cluster

haclus -value *attribute* [*cluster*] [-localclus]

The attribute must be specified in this command. If you do not specify the cluster name, the command displays the attribute value on the local cluster.

To display the state of a local or remote cluster

haclus -state [*cluster*] [-localclus]

The variable *cluster* represents the cluster. If a cluster is not specified, the state of the local cluster and the state of all remote cluster objects as seen by the local cluster are displayed.

For information on the state of a local or remote cluster

haclus -display [*cluster*] [-localclus]

If a cluster is not specified, information on the local cluster is displayed.

For a list of local and remote clusters

haclus -list [*conditionals*]

Lists the clusters that meet the specified conditions, beginning with the local cluster.

To display usage for the cluster command

haclus [-help [-modify]]

To display the status of a faulted cluster

haclus -status *cluster*

Displays the status on the specified faulted cluster. If no cluster is specified, the command displays the status on all faulted clusters. It lists the service groups that were not in the OFFLINE or the FAULTED state before the fault occurred. It also suggests corrective action for the listed clusters and service groups.

Querying status

For the status of local and remote clusters

hastatus

Querying heartbeats

The `hahb` command is used to manage WAN heartbeats that emanate from the local cluster. Administrators can monitor the "health of the remote cluster via heartbeat commands and mechanisms such as Internet, satellites, or storage replication technologies. Heartbeat commands are applicable only on the cluster from which they are issued.

Note: You must have Cluster Administrator privileges to add, delete, and modify heartbeats.

The following commands are issued from the command line.

For a list of heartbeats configured on the local cluster

`hahb -list [`*`conditionals`*`]`

The variable *conditionals* represents the conditions that must be met for the heartbeat to be listed.

To display information on heartbeats configured in the local cluster

`hahb -display [`*`heartbeat ...`*`]`

If *heartbeat* is not specified, information regarding all heartbeats configured on the local cluster is displayed.

To display the state of the heartbeats in remote clusters

`hahb -state [`*`heartbeat`*`] [-clus `*`cluster`*`]`

For example, to get the state of heartbeat ICMP from the local cluster to the remote cluster phoenix:

```
hahb -state ICMP -clus phoenix
```

To display an attribute value of a configured heartbeat

hahb -value *heartbeat attribute* [-clus *cluster*]

The -value option provides the value of a single attribute for a specific heartbeat. The cluster name must be specified for cluster-specific attribute values, but not for global.

For example, to display the value of the ClusterList attribute for heartbeat ICMP:

```
hahb -value Icmp ClusterList
```

Note that ClusterList is a global attribute.

To display usage for the command hahb

hahb [-help [-modify]]

If the -modify option is specified, the usage for the hahb -modify option is displayed.

Administering global service groups

Operations for the VCS global clusters option are enabled or restricted depending on the permissions with which you log on. The privileges associated with each user role are enforced for cross-cluster, service group operations.

See "User privileges in global clusters" on page 66.

To bring a service group online across clusters for the first time

hagrp -online -force

To bring a service group online across clusters

hagrp -online *service_group* -sys *system* [-clus *cluster* | -localclus]

The option `-clus` brings the service group online on the system designated in the cluster. If a system is not specified, the service group is brought online on any node within the cluster. The option `-localclus` brings the service group online in the local cluster.

To bring a service group online on any node

```
hagrp -online [-force] service_group -any [-clus cluster | -
localclus]
```

The option `-any` specifies that HAD brings a failover group online on the optimal system, based on the requirements of service group workload management and existing group dependencies. If bringing a parallel group online, HAD brings the group online on each system designated in the SystemList attribute.

To take a service group offline across clusters

```
hagrp -offline [-force] [-ifprobed] service_group -sys system
[-clus cluster -localclus]
```

The option `-clus` takes offline the service group on the system designated in the cluster.

To take a service group offline anywhere

```
hagrp -offline [-ifprobed] service_group -any [-clus cluster |
       -localclus]
```

The option `-any` specifies that HAD takes a failover group offline on the system on which it is online. For a parallel group, HAD takes the group offline on each system on which the group is online. HAD adheres to the existing group dependencies when taking groups offline.

To switch a service group across clusters

```
hagrp -switch service_group -to system [-clus cluster
       -localclus]
```

The option `-clus` identifies the cluster to which the service group will be switched. The service group is brought online on the system specified by the -to system argument. If a system is not specified, the service group may be switched to any node within the specified cluster.

To switch a service group anywhere

```
hagrp -switch service_group -clus cluster
```

The option `-clus` identifies the cluster to which the service group will be switched. HAD then selects the target system on which to switch the service group.

Administering resources

To take action on a resource across clusters

```
hares -action resource token [-actionargs arg1 ...] [-sys system]
        [-clus cluster |-localclus]
```

The option `-clus` implies resources on the cluster. If the designated system is not part of the local cluster, an error is displayed. If the `-sys` option is not used, it implies resources on the local node.

To invoke the Info entry point across clusters

```
hares -refreshinfo resource [-sys system] [-clus cluster
        -localclus]
```

Causes the Info entry point to update the value of the ResourceInfo resource level attribute for the specified resource if the resource is online. If no system or remote cluster is specified, the Info entry point runs on local system(s) where the resource is online.

To display usage for the resource command

To display usage for the command `hares` and its various options:

```
hares [-help [-modify |-list]]
```

Administering clusters in global clusters

To add a remote cluster object

haclus -add *cluster ip*

The variable *cluster* represents the cluster. This command does not apply to the local cluster.

To delete a remote cluster object

haclus -delete *cluster*

The variable *cluster* represents the cluster.

To modify an attribute of a local or remote cluster object

haclus -modify *attribute value* [-clus *cluster*]*...*

The variable *cluster* represents the cluster.

To declare the state of a cluster after a disaster

haclus -declare *disconnet/outage/disaster/replica* -clus *cluster* [-failover]

The variable *cluster* represents the remote cluster.

Changing the cluster name

This section describes how to change the ClusterName in a global cluster configuration. The instructions describe how to rename VCSPriCluster to VCSPriCluster2 in a two-cluster configuration, comprising clusters VCSPriCluster and VCSSecCluster configured with the global group AppGroup.

Before changing the cluster name, make sure the cluster is not part of any ClusterList, in the wide-area Heartbeat agent and in global service groups.

To change the name of a cluster

1 Run the following commands from cluster VCSPriCluster:

```
hagrp -offline ClusterService -any
hagrp -modify AppGroup ClusterList -delete VCSPriCluster
haclus -modify ClusterName VCSPriCluster2
hagrp -modify AppGroup ClusterList -add VCSPriCluster2 0
```

2 Run the following commands from cluster VCSSecCluster:

```
hagrp -offline ClusterService -any
hagrp -modify appgrp ClusterList -delete VCSPriCluster
hahb -modify Icmp ClusterList -delete VCSPriCluster
haclus -delete VCSPriCluster
haclus -add VCSPriCluster2 your_ip_address
hahb -modify Icmp ClusterList -add VCSPriCluster2
hahb -modify Icmp Arguments your_ip_address -clus
   VCSPriCluster2
hagrp -modify AppGroup ClusterList -add VCSPriCluster2 0
hagrp -online ClusterService -any
```

3 Run the following command from the cluster renamed to VCSPriCluster2:

```
hagrp -online ClusterService -any
```

Administering heartbeats

To create a heartbeat

hahb -add *heartbeat*

For example, type the following command to add a new IcmpS heartbeat. This represents a heartbeat sent from the local cluster and immediately forks off the specified agent process on the local cluster.

```
hahb -add IcmpS
```

To modify a heartbeat

hahb -modify *heartbeat attribute value* ... [-clus *cluster*]

If the attribute is local, that is, it has a separate value for each remote cluster in the ClusterList attribute, the option `-clus` *cluster* must be specified. Use `-delete -keys` to clear the value of any list attributes.

For example, type the following command to modify the ClusterList attribute and specify targets “phoenix and “houston for the newly created heartbeat:

```
hahb -modify ICMP ClusterList phoenix houston
```

To modify the Arguments attribute for target phoenix:

```
hahb -modify ICMP Arguments phoenix.veritas.com
       -clus phoenix
```

To delete a heartbeat

hahb -delete *heartbeat*

To change the scope of an attribute to cluster-specific

hahb -local *heartbeat attribute*

For example, type the following command to change the scope of the attribute AYAInterval from global to cluster-specific:

```
hahb -local ICMP AYAInterval
```

To change the scope of an attribute to global

hahb -global *heartbeat attribute value ...*
| *key* ... | *key value* ...

For example, type the following command to change the scope of the attribute AYAInterval from cluster-specific to cluster-generic:

```
hahb -global ICMP AYAInterval 60
```

Chapter 21

Setting up replicated data clusters

- About replicated data clusters
- How VCS replicated data clusters work

About replicated data clusters

The Replicated Data Cluster (RDC) 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.)

A Replicated Data Cluster (RDC) uses data replication to assure data access to nodes. An RDC exists within a single VCS cluster. In an RDC configuration, if an application or a system fails, the application is failed over to another system within the current primary site. If the entire primary site fails, the application is migrated to a system in the remote secondary site (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 sites. The replication service group must be online at both sites simultaneously, and must be configured as a hybrid VCS service group.

The application service group is configured as a failover service group. The application 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 site and the disaster recovery secondary site 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 site.

Note: You must use dual dedicated LLT links between the replicated nodes.

How VCS replicated data clusters work

To understand how a replicated data cluster configuration works, let us take the example of an application configured in a VCS replicated data cluster. The configuration has two system zones:

- Primary zone (zone 0) comprising nodes located at the primary site and attached to the primary storage
- Secondary zone (zone 1) comprising nodes located at the secondary site and attached to the secondary storage

The application is installed and configured on all nodes in the cluster. Application data is located on shared disks within each RDC zone and is replicated across RDC zones to ensure data concurrency. The application service group is online on a system in the current primary zone and is configured to fail over in the cluster.

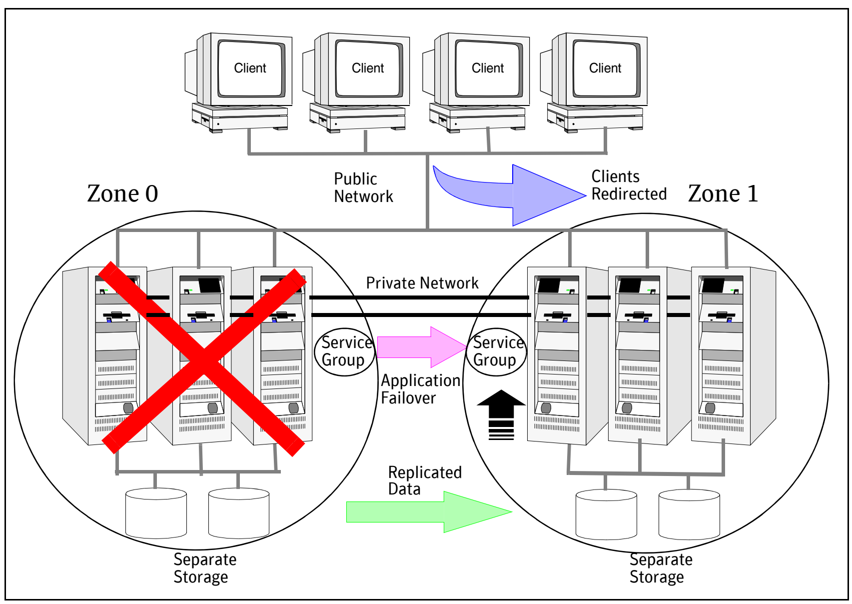

In the event of a system or application failure, VCS attempts to fail over the application service group to another system within the same RDC zone. However, in the event that VCS fails to find a failover target node within the primary RDC zone, VCS switches the service group to a node in the current secondary RDC zone (zone 1). VCS also redirects clients once the application is online on the new location.

Setting up a replicated data cluster configuration

See the *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL* for detailed configuration information.

Section VI

Troubleshooting and performance

- Chapter 22, "VCS performance considerations" on page 621
- Chapter 22, "Troubleshooting and recovery for VCS" on page 601

Chapter 22

VCS performance considerations

- How cluster components affect performance
- How cluster operations affect performance
- Monitoring CPU usage
- VCS agent statistics

How cluster components affect performance

VCS and its agents run on the same systems as the applications. Therefore, VCS attempts to minimize its impact on overall system performance. The three main components of clustering that have an impact on performance include the kernel; specifically, GAB and LLT, the VCS engine (HAD), and the VCS agents. For details on attributes or commands mentioned in the following sections, see the chapter on administering VCS from the command line and the appendix on VCS attributes.

Kernel components (GAB and LLT)

Typically, overhead of VCS kernel components is minimal. Kernel components provide heartbeat and atomic information exchange among cluster systems. By default, each system in the cluster sends two small heartbeat packets per second to other systems in the cluster. Heartbeat packets are sent over all network links configured in the /etc/llttab configuration file. System-to-system communication is load-balanced across all private network links. If a link fails, VCS continues to use all remaining links. Typically, network links are private and do not increase traffic on the public network or LAN. You can configure a public network (LAN) link as low-priority, which by default generates a small (approximately 64-byte) broadcast packet per second from each system, and which will carry data only when all private network links have failed.

The VCS engine (HAD)

The VCS engine, HAD, runs as a daemon process. By default it runs as a high-priority process, which ensures it sends heartbeats to kernel components and responds quickly to failures. HAD runs logging activities in a separate thread to reduce the performance impact on the engine due to logging.

VCS waits in a loop waiting for messages from agents, `ha` commands, the graphical user interfaces, and the other systems. Under normal conditions, the number of messages processed by HAD is few. They mainly include heartbeat messages from agents and update messages from the global counter. VCS may exchange additional messages when an event occurs, but typically overhead is nominal even during events. Note that this depends on the type of event; for example, a resource fault may involve taking the group offline on one system and bringing it online on another system. A system fault invokes failing over all online service groups on the faulted system.

To continuously monitor VCS status, use the VCS graphical user interfaces or the command `hastatus`. Both methods maintain connection to VCS and register for events, and are more efficient compared to running commands like `hastatus -summary` or `hasys` in a loop.

The number of clients connected to VCS can affect performance if several events occur simultaneously. For example, if five GUI processes are connected to VCS, VCS sends state updates to all five. Maintaining fewer client connections to VCS reduces this overhead.

How agents impact performance

The VCS agent processes have the most impact on system performance. Each agent process has two components: the agent framework and the agent entry points. The agent framework provides common functionality, such as communication with the HAD, multithreading for multiple resources, scheduling threads, and invoking entry points. Agent entry points implement agent-specific functionality. Follow the performance guidelines below when configuring agents.

Monitoring resource type and agent configuration

By default, VCS monitors each resource every 60 seconds. You can change this by modifying the MonitorInterval attribute for the resource type. You may consider reducing monitor frequency for non-critical or resources with expensive monitor operations. Note that reducing monitor frequency also means that VCS may take longer to detect a resource fault.

By default, VCS also monitors offline resources. This ensures that if someone brings the resource online outside of VCS control, VCS detects it and flags a concurrency violation for failover groups. To reduce the monitoring frequency of offline resources, modify the OfflineMonitorInterval attribute for the resource type.

The VCS agent framework uses multithreading to allow multiple resource operations to run in parallel for the same type of resources. For example, a single Mount agent handles all mount resources. The number of agent threads for most resource types is 10 by default. To change the default, modify the NumThreads attribute for the resource type. The maximum value of the NumThreads attribute is 30.

Continuing with this example, the Mount agent schedules the `monitor` entry point for all mount resources, based on the MonitorInterval or OfflineMonitorInterval attributes. If the number of mount resources is more than NumThreads, the monitor operation for some mount resources may be required to wait to execute the `monitor` entry point until the thread becomes free.

Additional considerations for modifying the NumThreads attribute include:

- If you have only one or two resources of a given type, you can set NumThreads to a lower value.
- If you have many resources of a given type, evaluate the time it takes for the `monitor` entry point to execute and the available CPU power for monitoring. For example, if you have 50 mount points, you may want to increase NumThreads to get the ideal performance for the Mount agent without affecting overall system performance.

You can also adjust how often VCS monitors various entry points by modifying their associated attributes. The attributes MonitorTimeout, OnlineTimeOut, and OfflineTimeout indicate the maximum time (in seconds) within which the monitor, online, and offline entry points must complete or else be terminated. The default for the MonitorTimeout attribute is 60 seconds. The defaults for the OnlineTimeout and OfflineTimeout attributes is 300 seconds. For best results, Symantec recommends measuring the time it takes to bring a resource online, take it offline, and monitor before modifying the defaults. Issue an online or offline command to measure the time it takes for each action. To measure how long it takes to monitor a resource, fault the resource and issue a probe, or bring the resource online outside of VCS control and issue a probe.

Agents typically run with normal priority. When you develop agents, consider the following:

- If you write a custom agent, write the monitor entry point using C or C++. If you write a script-based monitor, VCS must invoke a new process each time with the monitor. This can be costly if you have many resources of that type.

- If monitoring the resources is proving costly, you can divide it into cursory, or shallow monitoring, and the more extensive deep (or in-depth) monitoring. Whether to use shallow or deep monitoring depends on your configuration requirements.

Additional considerations for agents

Properly configure the attribute SystemList for your service group. For example, if you know that a service group can go online on sysa and sysb only, *do not* include other systems in the SystemList. This saves additional agent processes and monitoring overhead.

The VCS graphical user interfaces

The VCS graphical user interfaces, Cluster Manager (Java Console) and Cluster Management Console maintain a persistent connection to HAD, from which they receive regular updates regarding cluster status. For best results, run the Java and Web Consoles on a system outside the cluster to avoid impact on node performance.

How cluster operations affect performance

This section describes how operations on systems, resources, and service groups in the cluster affect performance.

Booting a cluster system

When a cluster system boots, the kernel drivers and VCS process start in a particular order. If it is the first system in the cluster, VCS reads the cluster configuration file main.cf and builds an "in-memory configuration database. This is the LOCAL_BUILD state. After building the configuration database, the system transitions into the RUNNING mode. If another system joins the cluster while the first system is in the LOCAL_BUILD state, it must wait until the first system transitions into RUNNING mode. The time it takes to build the configuration depends on the number of service groups in the configuration and their dependencies, and the number of resources per group and resource dependencies. VCS creates an object for each system, service group, type, and resource. Typically, the number of systems, service groups and types are few, so the number of resources and resource dependencies determine how long it takes to build the configuration database and get VCS into RUNNING mode. If a system joins a cluster in which at least one system is in RUNNING mode, it builds the configuration from the lowest-numbered system in that mode.

Note: Bringing service groups online as part of AutoStart occurs after VCS transitions to RUNNING mode.

When a resource comes online

The online entry point of an agent brings the resource online. This entry point may return before the resource is fully online. The subsequent monitor determines if the resource is online, then reports that information to VCS. The time it takes to bring a resource online equals the time for the resource to go online, plus the time for the subsequent monitor to execute and report to VCS.

Most resources are online when the online entry point finishes. The agent schedules the monitor immediately after the entry point finishes, so the first monitor detects the resource as online. However, for some resources, such as a database server, recovery can take longer. In this case, the time it takes to bring a resource online depends on the amount of data to recover. It may take multiple monitor intervals before a database server is reported online. When this occurs, it is important to have the correct values configured for the OnlineTimeout and OnlineWaitLimit attributes of the database server resource type.

When a resource goes offline

Similar to the online entry point, the offline entry point takes the resource offline and may return before the resource is actually offline. Subsequent monitoring confirms whether the resource is offline. The time it takes to offline a resource equals the time it takes for the resource to go offline, plus the duration of subsequent monitoring and reporting to VCS that the resource is offline. Most resources are typically offline when the offline entry point finishes. The agent schedules the monitor immediately after the offline entry point finishes, so the first monitor detects the resource as offline.

When a service group comes online

The time it takes to bring a service group online depends on the number of resources in the service group, the service group dependency structure, and the time to bring the group's resources online. For example, if service group G1 has three resources, R1, R2, and R3 (where R1 depends on R2 and R2 depends on R3), VCS first onlines R3. When R3 is online, VCS onlines R2. When R2 is online, VCS onlines R1. The time it takes to online G1 equals the time it takes to bring all resources online. However, if R1 depends on both R2 and R3, but there was no dependency between them, the online operation of R2 and R3 is started in parallel. When both are online, R1 is brought online. The time it takes to online the group is Max (the time to online R2 and R3), plus the time to online R1. Typically, broader service group trees allow more parallel operations and can be brought online faster. More complex service group trees do not allow much parallelism and serializes the group online operation.

The time it takes to bring a service group online or take it offline also depends on the type of service group, such as fileshare, printshare, enterprise agent, etc.

For a fileshare service group, there are four factors that determine how long it takes to bring a fileshare online:

- ShareSubDirectories
  If set to 1, each child subdirectory is shared. the fileshare group's online entry point shares child folders in addition to parent folders.
- Number of subdirectories
  The greater the number of subdirectories being shared, the longer it takes to bring online, monitor, and take offline a fileshare service group.
- Number of permissions
  For each share, the online entry point applies the share permissions as configured.
- AutoShare and AutoControl

By default, if ShareSubDirectories is set, the fileshare service group monitors new directories and shares them. AutoShare occurs in the monitor entry points.

For a printshare service group, the number of printers configured in the service group determines the time required for the service group to come online. The greater the number of printers, the more time required to bring the group online, monitor it, and take it offline.

When a service group goes offline

Taking service groups offline works from the top down, as opposed to the online operation, which works from the bottom up. The time it takes to offline a service group depends on the number of resources in the service group and the time to offline the group's resources. For example, if service group G1 has three resources, R1, R2, and R3, VCS first offlines R1. When R1 is offline, VCS offlines R2. When R2 is offline, VCS offlines R3. The time it takes to offline G1 equals the time it takes for all resources to go offline.

When a resource fails

The time it takes to detect a resource fault or failure depends on the MonitorInterval attribute for the resource type. When a resource faults, the next monitor detects it. The agent may not declare the resource as faulted if the ToleranceLimit attribute is set to non-zero. If the `monitor` entry point reports offline more often than the number set in ToleranceLimit, the resource is declared faulted. However, if the resource remains online for the interval designated in the ConfInterval attribute, previous reports of offline are not counted against ToleranceLimit.

When the agent determines that the resource is faulted, it calls the clean entry point (if implemented) to verify that the resource is completely offline. The monitor following clean verifies the offline. The agent then tries to restart the resource according to the number set in the RestartLimit attribute (if the value of the attribute is non-zero) before it gives up and informs HAD that the resource is faulted. However, if the resource remains online for the interval designated in ConfInterval, earlier attempts to restart are not counted against RestartLimit.

In most cases, ToleranceLimit is 0. The time it takes to detect a resource failure is the time it takes the agent monitor to detect failure, plus the time to clean up the resource if the clean entry point is implemented. Therefore, the time it takes to detect failure depends on the MonitorInterval, the efficiency of the monitor and clean (if implemented) entry points, and the ToleranceLimit (if set).

When a system fails

When a system crashes or is powered off, it stops sending heartbeats to other systems in the cluster. By default, other systems in the cluster wait 21 seconds before declaring it dead. The time of 21 seconds derives from 16 seconds default timeout value for LLT peer inactive timeout, plus 5 seconds default value for GAB stable timeout. The default peer inactive timeout is 16 seconds, and can be modified in the /etc/llttab file. For example, to specify 12 seconds:

```
set-timer peerinact:1200
```

Note: After modifying the peer inactive timeout, you must unconfigure, then restart LLT before the change is implemented. To unconfigure LLT, type `lltconfig -u`. To restart LLT, type `lltconfig -c`.

GAB stable timeout can be changed by specifying:

```
gabconfig -t timeout_value_milliseconds
```

Though this can be done, we *do not* recommend changing the values of the LLT peer inactive timeout and GAB stable timeout.

If a system reboots, it becomes unavailable until the reboot is complete. The reboot process kills all processes, including HAD. When the VCS process is killed, other systems in the cluster mark all service groups that can go online on the rebooted system as autodisabled. The AutoDisabled flag is cleared when the system goes offline. As long as the system goes offline within the interval specified in the ShutdownTimeout value, VCS treats this as a system reboot. The ShutdownTimeout default value of 120 can be changed by modifying the attribute.

See “System attributes” on page 728.

When a network link fails

If a system loses a network link to the cluster, other systems stop receiving heartbeats over the links from that system. As mentioned above, LLT detects this and waits for 16 seconds before declaring the system lost a link.

When a system panics

There are several instances in which GAB will intentionally panic a system, including if it detects an internal protocol error or discovers an LLT node-ID conflict. Three other instances are described below.

Client process failure

If a client process fails to heartbeat to GAB, the process is killed. If the process hangs in the kernel and cannot be killed, GAB halts the system. If the `-k` option is used in the `gabconfig` command, GAB tries to kill the client process until successful, which may have an impact on the entire cluster. If the `-b` option is used in `gabconfig`, GAB does not try to kill the client process. Instead, it panics the system when the client process fails to heartbeat. This option cannot be turned off once set.

HAD heartbeats with GAB at regular intervals. The heartbeat timeout is specified by HAD when it registers with GAB; the default is 15 seconds. If HAD gets stuck within the kernel and cannot heartbeat with GAB within the specified timeout, GAB tries to kill HAD by sending a SIGABRT signal. If it does not succeed, GAB sends a SIGKILL and closes the port. By default, GAB tries to kill HAD five times before closing the port. The number of times GAB tries to kill HAD is a kernel tunable parameter, gab_kill_ntries, and is configurable. The minimum value for this tunable is 3 and the maximum is 10.

This is an indication to other nodes that HAD on this node has been killed. Should HAD recover from its stuck state, it first processes pending signals. Here it will receive the SIGKILL first and get killed.

After sending a SIGKILL, GAB waits for a specific amount of time for HAD to get killed. If HAD survives beyond this time limit, GAB panics the system. This time limit is a kernel tunable parameter, *gab_isolate_time* and is configurable. The minimum value for this timer is 16 seconds and maximum is 4 minutes.

Network failure

If a network partition occurs, a cluster can “split into two or more separate sub-clusters. When two clusters join as one, VCS designates that one system be ejected. GAB prints diagnostic messages and sends iofence messages to the system being ejected. The system receiving the iofence messages tries to kill the client process. The `-k` option applied here. If the `-j` option is used in `gabconfig`, the system is halted when the iofence message is received.

Quick reopen

If a system leaves cluster and tries to join the cluster before the new cluster is configured (default is five seconds), the system is sent an iofence message with reason set to “quick reopen. When the system receives the message, it tries to kill the client process.

When a service group switches over

The time it takes to switch a service group equals the time to offline a service group on the source system, plus the time to bring the service group online on the target system.

When a service group fails over

The time it takes to fail over a service group when a resource faults equals

- the time it takes to detect the resource fault
- the time it takes to offline the service group on source system
- the time it takes for the VCS policy module to select target system
- the time it takes to bring the service group online on target system

The time it takes to fail over a service group when a system faults equals

- the time it takes to detect system fault
- the time it takes to offline the service group on source system
- the time it takes for the VCS policy module to select target system
- the time it takes to bring the service group online on target system

The time it takes the VCS policy module to determine the target system is negligible in comparison to the other factors.

If you have a firm group dependency and the child group faults, VCS offlines all immediate and non-immediate parent groups before bringing the child group online on the target system. Therefore, the time it takes a parent group to be brought online also depends on the time it takes the child group to be brought online.

Monitoring CPU usage

VCS includes a system attribute, CPUUsageMonitoring, which monitors CPU usage on a specific system and notifies the administrator when usage has been exceeded.

The default values for the CPUUsageMonitoring attribute are:

- Enabled = 0
- NotifyThreshold = 0
- NotifyTimeLimit = 0
- ActionThreshold = 0
- ActionTimeLimit = 0
- Action = NONE.

The values for ActionTimeLimit and NotifyTimeLimit represent the time in seconds. The values for ActionThreshold and NotifyThreshold represent the threshold in terms of CPU percentage utilization.

If Enabled is set to 1, HAD monitors the usage and updates CPUUsage attribute. If Enabled is set to 0 (default), HAD does not monitor the usage.

If the system's CPU usage continuously exceeds the value set in NotifyThreshold for a duration greater than the value set in NotifyTimeLimit, HAD sends notification via an SNMP trap or SMTP message.

If the CPU usage continuously exceeds the value set in NotifyThreshold for a duration greater than the value set in NotifyTimeLimit, subsequent notifications are sent after five minutes to avoid sending notifications too frequently (if the NotifyTimeLimit value is set to a value less than five minutes). In this case, notification is sent after the first interval of NotifyTimeLimit. As CPU usage continues to exceed the threshold value, notifications are sent after five minutes. If the values of NotifyThreshold or NotifyTimeLimit are set to 0, no notification is sent.

If system's CPU usage exceeds the value set in ActionThreshold continuously for a duration greater than the value set in ActionTimeLimit, the specified action is taken. If the CPU usage continuously exceeds the ActionThreshold for a duration greater than the value set in ActionTimeLimit, subsequent action is taken after five minutes to avoid taking action too frequently (if the ActionTimeLimit value is set to less than five minutes). In this case action is taken after the first interval of ActionTimeLimit. As CPU usage continues to exceed the threshold value, action is taken after five minutes. If the values of ActionThreshold or ActionTimeLimit are set to 0, no action is taken. Actions can have one of the following values:

NONE: No action will be taken and the message is logged in the VCS engine log.

REBOOT: System is rebooted.

CUSTOM: The cpuusage trigger is invoked.

VCS agent statistics

You can configure VCS to track the time taken for monitoring resources.

You can also detect potential problems with resources and systems on which resources are online by analyzing the trends in the time taken by the resource's monitor cycle. Note that VCS keeps track of monitor cycle times for online resources only.

VCS calculates the time taken for a monitor cycle to complete and computes an average of monitor times after a specific number of monitor cycles and stores the average in a resource-level attribute.

VCS also tracks increasing trends in the monitor cycle times and sends notifications about sudden and gradual increases in monitor times.

VCS uses the following parameters to compute the average monitor time and to detect increasing trends in monitor cycle times:

- *Frequency:* The number of monitor cycles after which the monitor time average is computed and sent to the VCS engine.
  For example, if Frequency is set to 10, VCS computes the average monitor time after every 10 monitor cycles.
- *ExpectedValue:* The expected monitor time (in milliseconds) for a resource.
  VCS sends a notification if the actual monitor time exceeds the expected monitor time by the ValueThreshold. So, if you set this attribute to 5000 for a FileOnOff resource, and if ValueThreshold is set to 40%, VCS will send a notification only when the monitor cycle for the FileOnOff resource exceeds the expected time by over 40%, that is 7000 milliseconds.
- *ValueThreshold:* The maximum permissible deviation (in percent) from the expected monitor time. When the time for a monitor cycle exceeds this limit, VCS sends a notification about the sudden increase or decrease in monitor time.
  For example, a value of 100 means that VCS sends a notification if the actual monitor time deviates from the expected time by over 100%.
  VCS sends these notifications conservatively. If 12 consecutive monitor cycles exceed the threshold limit, VCS sends a notification for the first spike, and then a collective notification for the next 10 consecutive spikes.
- *AvgThreshold:* The threshold value (in percent) for increase in the average monitor cycle time for a resource.

VCS maintains a running average of the time taken by the monitor cycles of a resource. The first such computed running average is used as a benchmark average. If the current running average for a resource differs from the benchmark average by more than this threshold value, VCS regards this as a sign of gradual increase or decrease in monitor cycle times and sends a notification about it for the resource. Whenever such an event occurs, VCS resets the internally maintained benchmark average to this new average. VCS sends notifications regardless of whether the deviation is an increase or decrease in the monitor cycle time.

For example, a value of 25 means that if the actual average monitor time is 25% more than the benchmark monitor time average, VCS sends a notification.

Tracking monitor cycle times

VCS marks sudden changes in monitor times by comparing the time taken for each monitor cycle with the ExpectedValue. If this difference exceeds the ValueThreshold, VCS sends a notification about the sudden change in monitor time. Note that VCS sends this notification only if monitor time increases.

VCS marks gradual changes in monitor times by comparing the benchmark average and the moving average of monitor cycle times. VCS computes the benchmark average after a certain number of monitor cycles and computes the moving average after every monitor cycle. If the current moving average exceeds the benchmark average by more than the AvgThreshold, VCS sends a notification about this gradual change in the monitor cycle time.

VCS attributes enabling agent statistics

This section describes the attributes that enable VCS agent statistics.

MonitorStatsParam	A resource type-level attribute, which stores the required parameter values for calculating monitor time statistics. `static str MonitorStatsParam = { Frequency = 10, ExpectedValue = 3000, ValueThreshold = 100, AvgThreshold = 40 }` ■ *Frequency:* Defines the number of monitor cycles after which the average monitor cycle time should be computed and sent to the engine. If configured, the value for this attribute must be between 1 and 30. It is set to 0 by default. ■ *ExpectedValue:* The expected monitor time in milliseconds for all resources of this type. Default=3000. ■ *ValueThreshold:* The acceptable percentage difference between the expected monitor cycle time (ExpectedValue) and the actual monitor cycle time. Default=100. ■ *AvgThreshold:* The acceptable percentage difference between the benchmark average and the moving average of monitor cycle times. Default=40
MonitorTimeStats	Stores the average time taken by a number of monitor cycles specified by the Frequency attribute along with a timestamp value of when the average was computed. `str MonitorTimeStats{} = { Avg = "0", TS = "" }` This attribute is updated periodically after a number of monitor cycles specified by the Frequency attribute. If Frequency is set to 10, the attribute stores the average of 10 monitor cycle times and is updated after every 10 monitor cycles. The default value for this attribute is 0.
ComputeStats	A flag that specifies whether VCS keeps track of the monitor times for the resource. `bool ComputeStats = 0` The value 0 indicates that VCS will not keep track of the time taken by the monitor routine for the resource. The value 1 indicates that VCS keeps track of the monitor time for the resource. The default value for this attribute is 0.

VCS performance with non-HA products

To ensure optimum performance, it is important to evaluate the impact of non-HA products on cluster nodes. Evaluating factors such as the complexity of the VCS configuration, the capacity of the hardware to host multiple applications, and the intended use of the product will assist you in determining how and where to host the applications.

When modifying the system, consider whether or not the change will cause the service group to fault. A simple task such as Windows Explorer browsing fileshares hosted by VCS may seem harmless, but it would prevent VCS from failing over because the drive is locked by another application.

VCS performance with SFW

If you use Veritas Storage Foundation for Windows (SFW) on clustered nodes, we strongly recommend the following:

- Carefully evaluate changes to underlying storage. Typically, changes to the volume and disk group configurations require corresponding changes to the VCS configuration. Common changes include unassigning or reassigning the drive letters, splitting or joining a disk group, or snapshotting the volume. Prior to implementing these types of changes, evaluate your configuration to determine whether to freeze, offline, or fail over the VCS service groups to avoid faulting the groups inadvertently.
- Like Cluster Manager, the SFW GUI runs under the Java Runtime environment and maintains a persistent connection to the SFW engine, from which it receives regular updates regarding status. For best results, run the SFW GUI on a system outside the cluster. This will avoid potential impact on node performance.
- Certain SFW operations, such as rescan, resync, etc., are CPU-intensive and can affect VCS performance. The VCS kernel module GAB expects the VCS engine, HAD, to send heartbeats that ensure the engine is functioning properly. If the heartbeat interval exceeds five seconds the engine logs an error.

By default, if GAB does not receive a heartbeat from HAD within 15 seconds, GAB assumes something is wrong and kills HAD (which then gets restarted by hashadow). You can tune this interval by changing the value of the system variable VCS_GAB_TIMEOUT, which specifies the number of seconds GAB waits for a heartbeat before killing HAD.

Troubleshooting and recovery for VCS

This chapter explains VCS unified logging and defines the message format. This chapter also describes how to troubleshoot common problems.

Logging

VCS generates two error message logs: the engine log and the agent log. Log file names are appended by letters. Letter A indicates the first log file, B the second, C the third, and so on.

The engine log is located at %VCS_HOME%\log\engine_A.txt. The format of engine log messages is:

```
Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI| Message Text
```

- *Timestamp*: the date and time the message was generated.
- *Mnemonic*: the string ID that represents the product (for example, VCS).
- *Severity*: levels include CRITICAL, ERROR, WARNING, NOTICE, and INFO (most to least severe, respectively).
- *UMI*: a unique message ID.
- *Message Text*: the actual message generated by VCS.

A typical engine log resembles:

```
2003/02/10 16:08:09 VCS INFO V-16-1-10077 received new cluster
membership.
```

The agent log is located at %VCS_HOME%\log*agent*_A.txt. The format of agent log messages is:

```
Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI | Agent Type |
Resource Name | Entry Point | Message Text
```

A typical agent log resembles:

```
2003/02/23 10:38:23 VCS WARNING V-16-2-23331
Oracle:VRT:monitor:Open for ora_lgwr failed, setting cookie to
null.
```

VCW logs

The VCS Configuration Wizard (VCW) log is located at `%allusersprofile%\Application Data\Veritas\Cluster Server\vcw.log`.

Here, *`%allusersprofile%`* is the file system directory containing application data for all users. A typical path is `C:\Documents and Settings\All Users\`.

The format of the VCW log is

ThreadID | Message Text

- *ThreadID*: the ID of the thread initiated by VCW.
- *Message Text*: the actual message generated by VCW.

A typical VCW log resembles:

```
00000576-00000264: ExecMethod return 00000000.
00000576-00000110: CRegistry::Query for VCS License failed.
Error=0x00000000
00000576-00000264: ExecMethod return 00000000.
00000576-00000264: ExecMethod return 00000001.
00000576-00000127: QueryDWORDValue returned 0x00000001
00000576-00000132: CRegistry::Query for VxSS Root information
failed. Error=0x00000001
```

VCWsilent logs

The VCWsilent log is located at `<currentdirectory>\vcwsilent.log`.

Here, *`<currentdirectory>`* is the directory from where the VCWsilent.exe is run.

A typical VCWsilent log resembles:

```
00005540-00000064: 5540: STARTING - Discovering NICs on the
selected machines...
00009956-00000064: 9956: STARTING - Generating private network
related files...
00009956-00000048: 9956: COMPLETED - Gererating LLT host
files...
00009956-00000048: 9956: COMPLETED - Generating GAB tab files...
00009956-00000048: 9956: COMPLETED - Generating main.cf file...
00009956-00000064: 9956: STARTING - Configuring LLT on all the
nodes.
00009956-00000048: 9956: COMPLETED - Configuring LLT on all the
nodes.
```

Message catalogs

VCS includes multilingual support for message catalogs. Most binary message catalogs (BMCs), are stored in `%VCS_HOME%\messages\language\`. The catalogs gab.bmc and llt.bmc are stored in `%VCS_ROOT%\comms\messages\language\`. The variable *language* represents a two-letter abbreviation. For example, en represents English.

The VCS command-line interface displays error/success messages in any language supported by VCS. The hamsg command displays the VCS engine logs in VCS-supported languages.

The following table shows the complete list of BMCs.

Table 23-2 Binary message catalogs

Module Name	Description
`VRTSvcsAgfw.bmc`	VCS agent framework messages
`VRTSvcsAlerts.bmc`	Alerts messages
`VRTSvcsApi.bmc`	VCS API messages
`VRTSvcsCommon.bmc`	Common messages
`VRTSvcsHad.bmc`	VCS engine (HAD) messages
`VRTSvcsHbfw.bmc`	VCS heartbeat framework messages
`VRTSvcsTriggers.bmc`	VCS triggers messages
`VRTSvcsAgent`*platform*`.bmc`	VCS bundled agent messages
`VRTSvcs`*platformagent_name*`.bmc`	VCS enterprise agent messages
`VRTSvcsWac.bmc`	Wide-area connector messages
`gab.bmc`	GAB command-line interface messages
llt.bmc	LLt command-line interface messages

Handling network failure

VCS protects against network partitions by requiring that all systems be connected by two or more communication channels. In a VCS cluster, all systems send heartbeats to each other across communication channels. If a system's heartbeats are not received across one channel, VCS detects that the channel has failed. If a system's heartbeats are not received across any channels, VCS detects that the system has failed. The services running on that system are then restarted on another.

VCS continues to operate as a single cluster when at least one network channel exists between the systems. However, when only one channel remains, failover due to system failure is disabled. Even after the last network connection is lost, VCS continues to operate as partitioned clusters on each side of the failure. For more information on protecting your cluster against network failure, see"Verifying LLT, GAB, and cluster operation" on page 649.

Disabling failover

When VCS loses communication with a system, a new regular membership is issued that excludes the departed system. VCS must then determine if it should restart that system's services, or if the system is running services outside of communication with VCS. Two conditions indicate that the system could still be running the services:

- Prior to the system's departure, the systems remaining in the new membership were connected to the departed system by only one communication channel.
- The departed system continues writing heartbeats to disk. VCS detects these conditions using the jeopardy membership.

If there is at least one system in the new regular membership that was not part of the prior jeopardy membership, then failover is disabled only for those systems that left the regular membership and were part of the prior jeopardy membership. Failover is also disabled for systems that are in the new jeopardy membership and outside of the new regular membership. This indicates these systems are actively writing heartbeats to disk. If there are no systems in the new regular membership that were not part of the previous jeopardy membership, failover is disabled for all systems that have departed. This indicates that connections from the remaining systems to all systems in the prior regular membership were potentially unreliable.

Example of how VCS handles network failure

In the following example, a single cluster has two networks connecting four nodes.

Figure 23-2 VCS and network failure: Four node cluster

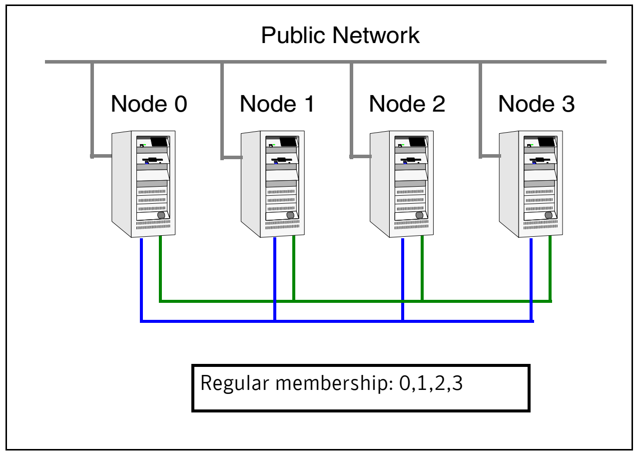

Jeopardy scenario: link failure

In this scenario, a link to node 2 fails, leaving the node with only one possible heartbeat.

Figure 23-3 VCS and network failure: Link to node 2 fails.

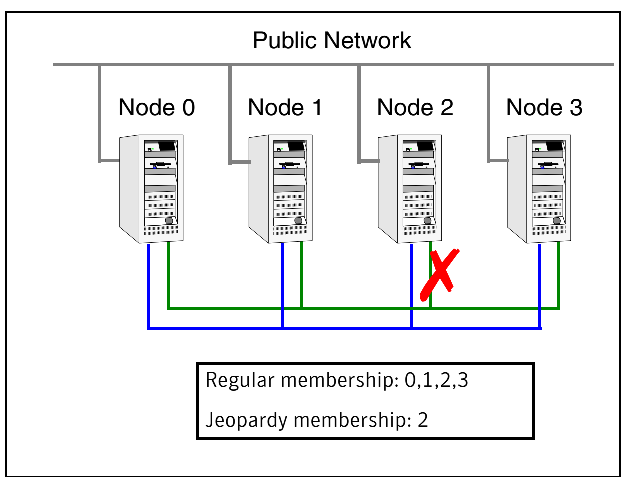

A new cluster membership is issued with nodes 0, 1, 2, and 3 in the regular membership and node 2 in a jeopardy membership. All normal cluster operations continue, including normal failover of service groups due to resource fault.

Jeopardy scenario: link and node failure

Consider that in the previous link-failure scenario, node 2 fails due to a power fault.

Figure 23-4 .VCS and network failure: Node 2 in jeopardy membership

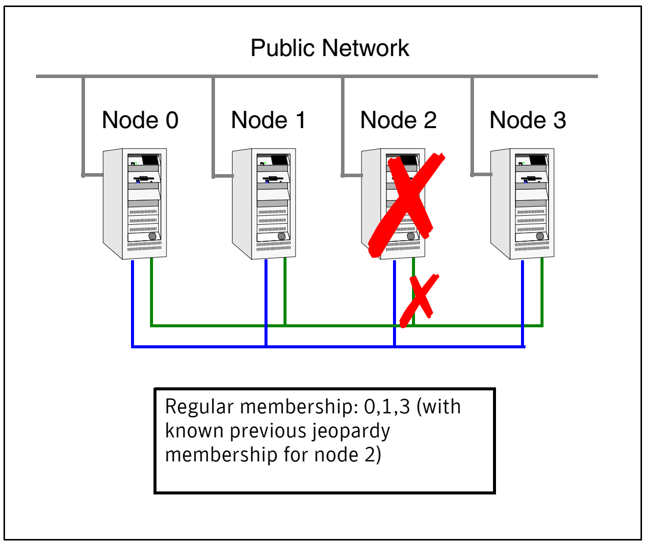

All other systems recognize that node 2 has faulted. In this situation, a new membership is issued for nodes 0, 1 and 3 as regular members. Since node 2 was in a jeopardy membership, service groups running on node 2 are autodisabled, so no other node can assume ownership of these service groups. If the node is actually failed, the system administrator can clear the AutoDisabled flag on the service groups in question and online the groups on other systems in the cluster.

Jeopardy scenario: failure of all links

In the scenario depicted in the illustration below, node 2 loses both heartbeats.

Figure 23-5 .VCS and network failure: Node 2 forms a single-node-mini cluster

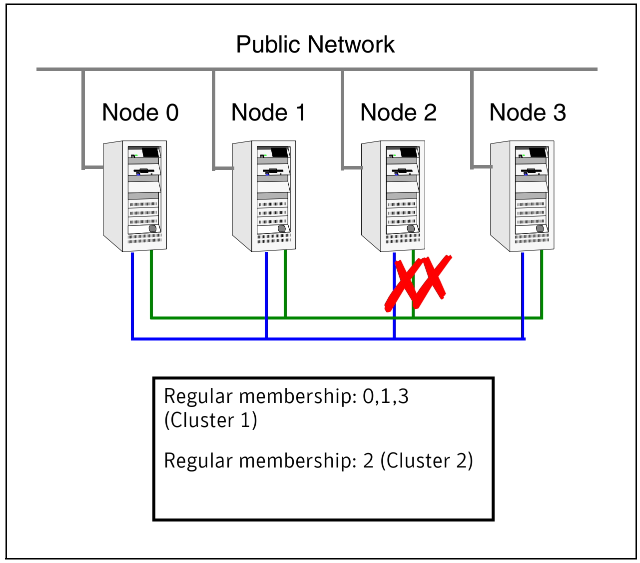

In this situation, a new membership is issued for node 0, 1 and 3 as regular members. Since node 2 was in a jeopardy membership, service groups running on node 2 are autodisabled, so no other node can assume ownership of these service groups. Nodes 0, 1 and 3 form a mini-cluster. Node 2 forms another single-node mini-cluster. All service groups that were present on nodes 0, 1 and 3 are autodisabled on node 2.

Network partitioning

With VCS, two or more communication channels guard against network partitioning; a condition where a failure on the network is misinterpreted as a failure of one or more systems in the cluster. If one system in the cluster assumes wrongly that another system has failed, it may restart applications already running on the other system, thereby corrupting the data.

Using a second communication channel enables VCS to distinguish between network and system failures. If all but one network channel fails, VCS enters a degraded mode that disables automatic application failover caused by system failure. If the last network channel fails, VCS splits into multiple “mini-clusters” without failing over or shutting down applications. This design enables administrative services to operate uninterrupted; for example, you can use VCS to shut down applications during system maintenance. When connections are restored, systems will attempt to rejoin into a single cluster. By default, GAB kills processes associated with ports on rejoining systems. To avoid potential data corruption during rejoin, add the option -j to the gabconfig command to enable system halt after a split. The gabconfig command is located in %VCS_ROOT\comms\gab.

When VCS shuts down a system

In some cases, VCS kernel components may intentionally bring down a system to avoid network partitioning. See the Veritas Cluster Server Release Notes for details.

Preexisting network partitions

A preexisting network partition refers to failures in communication channels that occur while the systems are down. Regardless of whether the cause is scheduled maintenance or system failure, VCS cannot respond to failures when systems are down. This leaves VCS vulnerable to network partitioning when the systems are booted. VCS seeding is designed to help prevent this situation.

VCS seeding

To protect your cluster from a pre-existing network partition, VCS employs the concept of a seed. Systems can be seeded automatically or manually. Note that only systems that have been seeded can run VCS.

By default, when a system comes up, it is not seeded. When the last system in a cluster is booted, the cluster will seed and start VCS on all systems. Systems can then be brought down and restarted in any combination. Seeding is automatic as long as at least one instance of VCS is running in the cluster.

Systems are seeded automatically in one of two ways:

- When an unseeded system communicates with a seeded system.
- When all systems in the cluster are unseeded and able to communicate with each other.

VCS requires that you declare the number of systems that will participate in the cluster.

When the last system is booted, the cluster will seed and start VCS on all systems. Systems can then be brought down and restarted in any combination. Seeding is automatic as long as at least one instance of VCS is running somewhere in the cluster.

Note that before VCS can accept HA commands, the cluster nodes must be seeded. If the nodes are not seeded and you attempt to issue a command, you receive the error message:

```
VCS:11037:Node has not received cluster membership yet, cannot
process HA command
```

To seed a cluster

1. Verify the value of gabconfig -c in the file %VCS_ROOT%\comms\gab\gabtab.txt is the same for all nodes.
2. Determine how many nodes are operational.
3. For each cluster node, modify gabtab.txt to reflect the required number of members to seed are equal to the number of cluster nodes in operation.
4. Reboot each node, or stop HAD -force on all nodes and restart.

Editing gabtab.txt after installing on a two-system cluster

In the unlikely event that the private network links fail, the -c -x option to the gabconfig command makes it possible to bring up VCS on a subset of cluster nodes. For an unseeded cluster, the -x option seeds the cluster, thus eliminating the requirement that all nodes are seeded before the cluster is brought up.

Note: Because the subset of nodes on which VCS was restarted maintains access to the disk and cannot communicate with the other cluster subset (which also has access to the disk), this option must be used with caution to avoid data corruption.

Reconnecting the private network

When a final network connection is lost, the systems on each side of the network partition segregate into mini-clusters.

Reconnecting a private network after a cluster has been segregated causes HAD to stop and restart. There are several rules that determine which systems will be affected.

- On a two-node cluster, the system with the lowest LLT host ID stays running and the higher recycles HAD.
- In a multi-node cluster, the largest running group stays running. The smaller groups recycle HAD.
- On a multi-node cluster splitting into two equal size clusters, the cluster with the lowest node number stays running. The higher group recycles HAD.

Troubleshooting VCS startup

When VCS is started, GAB, LLT, and HAD are started automatically. If they are not, review the corresponding log file. Startup errors for LLT and GAB are stored in the System Event log. Startup errors for HAD are stored in the Application Event log.

To view log files

1. From the Control Panel, double-click **Administrative Tools,** then **Event Viewer**.
2. Review the System Log to view LLT and GAB errors.
3. Review the Application Log to view HAD errors.

Low Latency Transport (LLT)

During installation, an llttab.txt configuration file containing minimum directives is created and placed in the following directory on each node in the cluster:

Drive:\Program Files\VERITAS\comms\llt

Each llttab.txt file specifies the node's ID, the network interfaces to use, and other directives. The most common LLT directives are listed on page 648.

Note: The directives must always appear as they are listed in the original default llttab.txt file.

Common LLT directives

link	Attaches LLT to a network interface. At least one link is required, and up to eight are supported. The first argument to `link` is a user-defined tag shown in the `lltstat` output to identify the link. It may also be used in llttab.txt to set optional static MAC addresses. The second argument to `link` is the device name of the network interface. (To obtain network device names, use the `objdir\device` command provided by the Windows 2000 Device Driver Kit.) There should be one `link` directive for each network interface, and each network interface configured for LLT must be attached to a separate physical network. LLT uses an unregistered Ethernet SAP of 0xCAFE. If the SAP is unacceptable, refer to the llttab.txt online Help for information on how to customize SAP. Note that IP addresses need not be assigned to the network device because LLT does not use IP addresses.
link-lowpri	Use this directive in place of `link` for public network interfaces. This directive prevents VCS communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. Note that LLT distributes network traffic evenly across all available network connections and broadcasts heartbeats to monitor each network connection.
set-cluster	Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. Note that LLT uses a default cluster number of zero.
set-node	Assigns the node ID. This number must be unique for each node in the cluster, and must be between 0-31. **Note:** LLT fails to operate if nodes share the same ID.
start	This directive must always appear last.

Group Membership Atomic Broadcast (GAB)

During installation, a gabtab.txt configuration file is automatically created and placed in the following directory on each system in the cluster:

Drive:\Program Files\VERITAS\comms\gab

Verifying LLT, GAB, and cluster operation

Before verifying LLT, GAB, or cluster operation, you must log on to any node in the cluster using an account with administrator privileges.

Verifying LLT

Use the `lltstat` command to verify the links are active for LLT. This command returns information about the LLT links for the node on which it is typed.

In the following example, `lltstat -n` is typed on System 0 and System 1 in a private network.

System 0

```
Drive:\> lltstat -n

LLT node information:

    Node                                  State      Links
    *0        HOUWIN201                   OPEN
2
     1        HOUWIN202                   OPEN
2
```

System 1

```
Drive:\> lltstat -n

LLT node information:

    Node                                  State      Links
     0        HOUWIN201                   OPEN
2
    *1        HOUWIN202                   OPEN
2
```

Note that each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which the command is typed.

If the output of `lltstat -n` does not show each node in the cluster, or does not show two links for each node, type `lltstat -nvv |` to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on System 0 in a private network. Note that each node should be OPEN, each link should be UP, and each address should be correct.

```
Drive:\> lltstat -nvv | more

Node    State       Link     Status   Address
*0      HOUWIN201   OPEN     Adapter0   UP  00:03:47:0D:A8:74
                             Adapter1   UP  00:03:47:0D:A8:75

1       HOUWIN202   OPEN     Adapter0   UP  00:03:47:0D:A4:46
                             Adapter1   UP  00:03:47:0D:A4:47

2                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN

3                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
4                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN

5                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN

6                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN

7                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN

8                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN

9                 CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
10                CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
12                CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
13                CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
14                CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
15                CONNWAIT   Adapter0   DOWN
                             Adapter1   DOWN
```

To obtain information about the ports open for LLT, type `lltstat -p` on any node. In the following example, `lltstat -p` is typed on System 0 in a private network.

```
Drive:\> lltstat -p

LLT port information:

    Port          Usage                          Cookie
```

```
    0              gab                        0x0
                   opens:                     0 1 2 3 4 5 6 7 8 9
10 11 12 13 14...
                   connects:                  0 1
```

Note that two nodes (0 and 1) are connected.

Verifying GAB

To verify GAB operation, type the following command as Administrator on each node:

Drive:\> **gabconfig -a**

If GAB is operating, the following GAB port membership information is returned:

```
    GAB Port Memberships
    ===================================
    Port           a       gen         a36e0003
membership 01
    Port           h       gen         fd570002
membership 01
```

Port a indicates GAB is communicating, gen a36e0003 is a random generation number, and membership 01 indicates nodes 0 and 1 are connected.

If GAB is not operating, no GAB port membership information is returned:

```
    GAB Port Memberships
    ===================================
```

If only one network is connected, the following GAB port membership information is returned:

```
    GAB Port Memberships
    ===================================
    Port           a       gen         a36e0003
membership                      01
    Port           a       gen         a36e0003
jeopardy     ;1
    Port           h       gen         fd570002
membership                      01
    Port           h       gen         fd570002
jeopardy     ;1
```

Verifying HAD

To verify HAD operation, type the following command as Administrator on each node:

Drive:\> **gabconfig -a**

If HAD is operating, the following port membership information is returned:

```
    GAB Port Memberships
    ===================================
    Port           a       gen        a36e0003
membership 01
    Port           h       gen        fd570002
membership 01
```

Port h indicates HAD is started, gen fd570002 is a random generation number, and membership 01 indicates nodes 0 and 1 are both running VCS.

If HAD is not operating, no port membership information is returned.

```
    GAB Port Memberships
    ===================================
```

See "VCS seeding" on page 645 for instructions on how to seed the cluster.

If HAD is running on only one node, the following port membership information is returned:

```
    GAB Port Memberships
    ===================================
    Port           a       gen        a36e0003       membership 01
    Port           h       gen        fd570002       membership 0
    Port           h       gen        fd51002        visible ;1
```

This information indicates HAD is running on node 1, but only GAB is running on node 0. Check the Application Event Log on node 0 for more information.

Verifying the cluster

To verify cluster operation, type the following command as Administrator on any node:

```
    Drive:\> hasys -display

#System            Attribute              Value
HOUWIN201          AgentsStopped          0
HOUWIN201          AvailableCapacity      100
HOUWIN201          CPUUsage               0
HOUWIN201          CPUUsageMonitoring     Enable 0 ActionThreshold 0
ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
HOUWIN201          Capacity               100
HOUWIN201          ConfigBlockCount       84
HOUWIN201          ConfigCheckSum         18907
HOUWIN201          ConfigDiskState        CURRENT
HOUWIN201          ConfigFile             C:\Program
Files\VERITAS\Cluster Server\conf\config
HOUWIN201          ConfigInfoCnt          0
HOUWIN201          ConfigModDate          Tue Dec 03 15:13:58 2001
```

```
HOUWIN201          CurrentLimits
HOUWIN201          DiskHbStatus
HOUWIN201          DynamicLoad            0
HOUWIN201          Frozen                 0
HOUWIN201          GUIIPAaddr
HOUWIN201          LLTNodeId              0
HOUWIN201          Limits
HOUWIN201          LinkHbStatus           Adapter0 UP Adapter1 UP
HOUWIN201          LoadTimeCounter        0
HOUWIN201          LoadTimeThreshold      600
HOUWIN201          LoadWarningLevel       80
HOUWIN201          MajorVersion           2
HOUWIN201          MinorVersion           0
HOUWIN201          NodeID                 0
HOUWIN201          OnGrpCnt               1
HOUWIN201          ShutdownTimeout        60
HOUWIN201          SourceFile             .\main.cf
HOUWIN201          SysInfo                WINNT:HOUWIN201,5.0,2195,
Service Pack 2, INTEL,1
HOUWIN201          SysName
HOUWIN201
HOUWIN201          SysState
RUNNING
HOUWIN201          SystemLocation
HOUWIN201          SystemOwner
HOUWIN201          TFrozen                0
HOUWIN201          TRSE                   0
HOUWIN201          UpDownState            Up
HOUWIN201          UserInt                0
HOUWIN201          UserStr                #
HOUWIN202          AgentsStopped          0
HOUWIN202          AvailableCapacity      100
HOUWIN202          CPUUsage               0
HOUWIN202          CPUUsageMonitoring     Enable 0 ActionThreshold 0
ActionTimeLimit 0 Action NONE NotifyThreshold 0NotifyTimeLimit 0
HOUWIN202          Capacity               100
HOUWIN202          ConfigBlockCount       84
HOUWIN202          ConfigCheckSum         18907
HOUWIN202          ConfigDiskState        CURRENT
HOUWIN202          ConfigFile             C:\Program Files\VERITAS\
Cluster Server\conf\config
HOUWIN202          ConfigInfoCnt          0
HOUWIN202          ConfigModDate          Tue Dec 03 15:15:58 2001
HOUWIN202          CurrentLimits
HOUWIN202          DiskHbStatus
HOUWIN202          DynamicLoad            0
HOUWIN202          Frozen                 0
HOUWIN202          GUIIPAaddr
HOUWIN202          LLTNodeIdHOUWIN202     Limits
HOUWIN202          LinkHbStatus           Adapter0 UP Adapter1 UP
HOUWIN202          LoadTimeCounter        0
HOUWIN202          LoadTimeThreshold      600
```

```
HOUWIN202          LoadWarningLevel       80
HOUWIN202          MajorVersion           2
HOUWIN202          MinorVersion           0
HOUWIN202          NodeID                 1
HOUWIN202          OnGrpCnt               1
HOUWIN202          ShutdownTimeout        60
HOUWIN202          SourceFile             .\main.cf
HOUWIN202          SysInfo                WINNT:HOUWIN202,5.0,2195,
Service Pack 2, INTEL,1
HOUWIN202          SysName
HOUWIN202
HOUWIN202          SysState               RUNNING
HOUWIN202          SystemLocation
HOUWIN202          SystemOwner
HOUWIN202          TFrozen                0
HOUWIN202          TRSE                   0
HOUWIN202          UpDownState            Up
HOUWIN202          UserInt                0
HOUWIN202          UserStr
```

Note the value for the attribute ConfigFile is an empty file created by default to enable VCS to start. Also note the value of the attribute SysState is RUNNING, which indicates VCS is started. This output indicates VCS was successfully installed on both nodes in the cluster.

VCS startup errors

This section includes error messages associated with starting VCS and provides descriptions of each error and the recommended action.

"VCS:10622 local configuration missing"

"VCS:10623 local configuration invalid"

"VCS:10624 local configuration stale"

The local configuration is invalid.

Recommended Action: Start the VCS engine, HAD, on another system that has a valid configuration file. The system with the configuration error "pulls" the valid configuration from the other system.

Another method is to correct the configuration file on the local system and force VCS to reread the configuration file. If the file appears valid, verify that is not an earlier version. It is possible that VCS marked the configuration stale by creating a .stale file because the last VCS shutdown was not graceful. The .stale file is created in the directory %VCS_HOME%\conf\config.

Type the following commands to verify the configuration and force VCS to reread the configuration file:

C:\> cd %VCS_HOME\conf\config

```
C:\> hacf -verify .
C:\> hasys -force system
```

"VCS:11032 registration failed. Exiting"

GAB was not registered or has become unregistered.

Recommended Action: GAB is registered by the `gabconfig` command in the file %VCS_ROOT%\comms\gab\gabtab.txt. Verify that the file exists and that it contains the command `gabconfig -c`.

GAB can become unregistered if LLT is set up incorrectly. Verify that the file is correct in %VCS_ROOT%\comms\llt\llttab.txt. If the LLT configuration is incorrect, make the appropriate changes and reboot.

"Waiting for cluster membership."

This indicates that GAB may not be seeded. If this is the case, the command `gabconfig -a` does not show any members, and the following messages may appear on the console or in the event log.

```
GAB: Port a registration waiting for seed port membership
GAB: Port h registration waiting for seed port membership
```

The following message will also be sent to the engine log:

```
Did not receive cluster membership, manual intervention may be
needed for seeding
```

Follow the instructions below to seed the cluster:

To seed the cluster

1. Verify the value of gabconfig -c in the file %VCS_ROOT%\comms\gab\gabtab.txt is the same for all nodes.
2. Determine how many nodes are operational.
3. For each cluster node, modify gabtab.txt to reflect the required number of members to seed are equal to the number of cluster nodes in operation.

Reboot each node, or stop HAD -force on all nodes and restart. See "VCS seeding" on page 645 for more information.

Troubleshooting secure clusters

"Error returned from engine: HAD on this node not accepting clients."

This error occurs when an HA command fails because the VCS engine could not initialize its security credentials. When this occurs, the following message is logged to the event log:

"Security ON. Init failed. Clients will be rejected."

Recommended Action:

- Verify the Symantec Product Authentication Service configuration. Make sure the cluster was configured to run in secure mode before the SecureClus attribute was set to 1. See "Enabling and disabling Symantec Product Authentication Service" on page 390 for instructions.
- Verify the Veritas Authentication Service is running. Stop and restart the service.
- Restart the VCS engine (HAD) on the node.

"Unable to connect to the VCS engine securely."

Recommended Action:

- Verify the Veritas Authentication Service is running. Stop and restart the service.
- Restart the VCS engine (HAD) on the node.

Troubleshooting service groups

This section cites the most common problems associated with bringing service groups online and taking them offline. Recommended action is also included, where applicable.

System is not in RUNNING state.

Recommended Action: Type `hasys -display` *system* to verify the system is running. See"System states" on page 696 for more information on system states.

Service group not configured to run on the system.

The SystemList attribute of the group may not contain the name of the system.

Recommended Action: Use the output of the command `hagrp -display` *service_group* to verify the system name.

Service group not configured to autostart.

If the service group is not starting automatically on the system, the group may not be configured to AutoStart, or may not be configured to AutoStart on that particular system.

Recommended Action: Use the output of the command `hagrp -display` *service_group* to verify the values of the AutoStart and AutoStartList attributes.

Service group is frozen.

Recommended Action: Use the output of the command `hagrp -display` *service_group* to verify the value of the Frozen and TFrozen attributes. Use the command `hagrp -unfreeze` to unfreeze the group. Note that VCS will not take a frozen service group offline.

Service group autodisabled.

When VCS does not know the status of a service group on a particular system, it autodisables the service group on that system. Autodisabling occurs under the following conditions:

- When the VCS engine, HAD, is not running on the system.
- When all resources within the service group are not probed on the system.
- When a particular system is visible through disk heartbeat only.

Under these conditions, all service groups that include the system in their SystemList attribute are autodisabled. *This does not apply to systems that are powered off.*

Recommended Action: Use the output of the command `hagrp -display` *service_group* to verify the value of the AutoDisabled attribute.

Caution: To bring a group online manually after VCS has autodisabled the group, make sure that the group is not fully or partially active on any system that has the AutoDisabled attribute set to 1 by VCS. Specifically, verify that all resources that may be corrupted by being active on multiple systems are brought down on the designated systems. Then, clear the AutoDisabled attribute for each system:

```
C:\> hagrp -autoenable service_group -sys system
```

Failover service group is online on another system.

The group is a failover group and is online or partially online on another system.

Recommended Action: Use the output of the command `hagrp -display` *service_group* to verify the value of the State attribute. Use the command `hagrp -offline` to offline the group on another system.

Service group is waiting for the resource to be brought online/taken offline.

Recommended Action: Review the IState attribute of all resources in the service group to locate which resource is waiting to go online (or which is waiting to be taken offline). Use the `hastatus` command to help identify the resource. See the engine and agent logs for information on why the resource is unable to be brought online or be taken offline.

To clear this state, make sure all resources waiting to go online/offline do not bring themselves online/offline. Use the command `hagrp -flush` to clear the internal state of VCS. You can now bring the service group online or take it offline on another system.

A critical resource faulted.

Output of the command `hagrp -display` *service_group* indicates that the service group has faulted.

Recommended Action: Use the command `hares -clear` to clear the fault.

Service group is waiting for a dependency to be met.

Recommended Action: To see which dependencies have not been met, type `hagrp -dep` *service_group* to view service group dependencies, or `hares -dep` *resource* to view resource dependencies.

Service group not fully probed.

This occurs if the agent processes have not monitored each resource in the service group. When the VCS engine, HAD, starts, it immediately "probes" to find the initial state of all of resources. (It cannot probe if the agent is not returning a value.) A service group must be probed on all systems included in the SystemList attribute before VCS attempts to bring the group online as part of AutoStart. This ensures that even if the service group was online prior to VCS being brought up, VCS will not inadvertently bring the service group online on another system.

Recommended Action: Use the output of `hagrp -display` *service_group* to see the value of the ProbesPending attribute for the system's service group. (It should be zero.) To determine which resources are not probed, verify the local Probed attribute for each resource on the specified system. Zero means waiting for probe result, 1 means probed, and 2 means VCS not booted. See the engine and agent logs for information.

ClusterService group configuration

If you run the `hastop -local` command on a node that is not defined in the ClusterService group's SystemList and has other service groups online, VCS takes the service groups offline on the node and the node gets stuck in the LEAVING state.

To prevent this from happening, make sure that the ClusterService group is defined on all nodes in the cluster. In other words, the SystemList attribute of the ClusterService group must contain all nodes in the cluster.

Troubleshooting resources

This section cites the most common problems associated with bringing resources online and taking them offline. Recommended action is also included, where applicable.

Service group brought online due to failover.

VCS attempts to bring resources online that were already online on the failed system, or were in the process of going online. Each parent resource must wait for its child resources to be brought online before starting.

Recommended Action: Verify that the child resources are online.

Waiting for service group states.

The state of the service group prevents VCS from bringing the resource online.

Recommended Action: See the appendix "Cluster and system states" for more information on states.

Waiting for child resources.

One or more child resources of parent resource are offline.

Recommended Action: Bring the child resources online first.

Waiting for parent resources.

One or more parent resources are online.

Recommended Action: Take the parent resources offline first.

Waiting for resource to respond.

The resource is waiting to come online or go offline, as indicated. VCS directed the agent to run an online entry point for the resource.

Recommended Action: Verify the resource's IState attribute. See the engine and agent logs for information on why the resource cannot be brought online.

Agent not running.

The resource's agent process is not running.

Recommended Action: Use `hastatus -summary` to see if the agent is listed as faulted. Restart the agent:

```
C:\> haagent -start resource_type -sys system
```

Invalid agent argument list.

The scripts are receiving incorrect arguments.

Recommended Action: Verify that the arguments to the scripts are correct. Use the output of `hares -display` *resource* to see the value of the ArgListValues attribute. If the ArgList attribute was dynamically changed, stop the agent and restart it.

To stop the agent, type:

```
C:\> haagent -stop resource_type -sys system
```

To restart the agent, type:

```
C:\> haagent -start resource_type -sys system
```

Troubleshooting notification

Occasionally you may encounter problems when using VCS notification. This section cites the most common problems and the recommended actions.

Notifier is configured but traps are not seen on SNMP console.

Recommended Action: Verify the version of SNMP traps supported by the console: VCS notifier sends SNMP v2.0 traps. If you are using HP OpenView Network Node Manager as the SNMP, verify events for VCS are configured using xnmevents. You may also try restarting the OpenView daemon (ovw) if, after merging VCS events in vcs_trapd, the events are not listed in the OpenView Network Node Manager Event configuration.

By default, notifier assumes the community string is public. If your SNMP console was configured with a different community, reconfigure it according to the notifier configuration.

Troubleshooting Cluster Management Console (Single Cluster Mode) or Web Console

Occasionally you may encounter problems when using Cluster Management Console (Single Cluster Mode) also referred to as Web Console. This section cites the most common problems and the recommended actions.

Unable to log on.

Recommended Action: Verify the user name exists and that the password is correct for the user name. Then verify your browser is Java, Javascript, and cookies enabled.

A Java NullPointerException occurs when connecting to a cluster from the Web Console.

Recommended Action: Verify the Security settings for the cluster. Make sure the VCS engine (had) and VCS Command Server both run in the same mode. For example, if the cluster is configured to run in secure mode, make sure HAD and VCS Command Server both run in secure mode. See"Enabling and disabling Symantec Product Authentication Service" on page 390 for more information.

Unable to view Cluster Manager on a browser using the Virtual IP/port number in URL (http://[virtual_ip:port_number]/vcs).

Recommended Action: Verify that the ClusterService service group, which has the IP and VRTSWebApp resources configured on it, is not offline or faulted on any node. If it is, use the command line to bring the group back online on at least one node.

Unable to view Cluster Manager on a browser using the HostName/port_number in URL (http://[host_name:port_number]/vcs.

Recommended Action: Verify that the host is running and that the ClusterService group is online on the host. If the host is down, access Cluster Manager (Web Console) using the URL http://*virtual_IP:port_number*/vcs. The cause of the failover should be apparent on Cluster Manager. Use Cluster Manager to administer nodes that are up and running in cluster.

Unable to bring the VCSweb resource online in the ClusterService group.

You cannot access the Web Console unless the VCSweb resource in the ClusterService group is online. The Web Console runs inside the VERITAS Web server (VRTSweb). VCSweb may fail to come online if the Web server cannot start because of one of the following reasons:

- ✔ Web server port unavailable: By default, the Web server binds itself to ports 8181, 8443, and 14300. If these ports are being used by another application, the Web server will fail to start.

  To determine if this is the reason, review the last few lines of the log file. If the output resembles the example below, the Web server port is already taken by another application:

```
5/28/03 8:13:35 PM PDT VRTSWEB INFO V-12-1-1041 Exception
encountered
LifecycleException: Protocol handler initialization failed:
java.net.BindException: Address already in use: JVM_Bind:8181
at  org.apache.coyote.tomcat4.CoyoteConnector.initialize
               (CoyoteConnector.java:1119)
at org.apache.catalina.startup.Embedded.start(Embedded.java:999)
at vrts.tomcat.server.VRTSweb.initServer(VRTSweb.java:2567)
at vrts.tomcat.server.VRTSweb.commandStartServer
               (VRTSweb.java:385
at vrts.tomcat.server.command.start.StartCommand.execute
               (StartCommand.java:59)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke
               (NativeMethodAccessorImpl.java:39)
at sun.reflect.DelegatingMethodAccessorImpl.invoke
               (DelegatingMethodAccessorImpl.java:25)
at java.lang.reflect.Method.invoke(Method.java:324)
at vrts.tomcat.bootstrap.Main.main(Main.java:243)
```

  Recommended Action: If you cannot make this port available for VRTSweb, refer to "Configuring ports for VRTSweb" on page 748 for instructions on how to change the value of the Web server port.

- ✔ Web server IP address unavailable: By default, the Web server binds itself to all IP addresses on the machine for the default ports 8181 and 8443. If you configure a specific IP address for the port, verify this IP address is available on the machine before the Web server starts. The Web server will fail to start if this IP address is not present on the machine.

  If the log output resembles the example below, the IP address is not available:

```
5/28/03 8:20:16 PM PDT VRTSWEB INFO V-12-1-1041 Exception
               encountered
LifecycleException: Protocol handler initialization failed:
java.net.BindException: Cannot assign requested address:
       JVM_Bind:8181
at org.apache.coyote.tomcat4.CoyoteConnector.initialize
               (CoyoteConnector.java:1119)
at org.apache.catalina.startup.Embedded.start(Embedded.java:999)
at vrts.tomcat.server.VRTSweb.initServer(VRTSweb.java:2567)
at vrts.tomcat.server.VRTSweb.commandStartServer
               (VRTSweb.java:385)
at vrts.tomcat.server.command.start.StartCommand.execute
```

```
                (StartCommand.java:59)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke
                (NativeMethodAccessorImpl.java:39)
at sun.reflect.DelegatingMethodAccessorImpl.invoke
                (DelegatingMethodAccessorImpl.java:25)
at java.lang.reflect.Method.invoke(Method.java:324)
at vrts.tomcat.bootstrap.Main.main(Main.java:243)
LifecycleException: Protocol handler initialization failed:
java.net.BindException: Cannot assign requested address:
                JVM_Bind:8181
at org.apache.coyote.tomcat4.CoyoteConnector.initialize
                (CoyoteConnector.java:1119)
at org.apache.catalina.startup.Embedded.start(Embedded.java:999)
at vrts.tomcat.server.VRTSweb.initServer(VRTSweb.java:2567)
at vrts.tomcat.server.VRTSweb.commandStartServer
                (VRTSweb.java:385)
at vrts.tomcat.server.command.start.StartCommand.execute
                (StartCommand.java:59)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke
            (NativeMethodAccessorImpl.java:39)
at sun.reflect.DelegatingMethodAccessorImpl.invoke
                (DelegatingMethodAccessorImpl.java:25)
at java.lang.reflect.Method.invoke(Method.java:324)
at vrts.tomcat.bootstrap.Main.main(Main.java:243)
```

Recommended Action: Make this IP address available on the machine and try to bring the VCSweb resource online again.

After reconfiguring virtual IP address, cannot access the Web Console using the new IP address.

Recommended Action: If ClusterService service group is online, changes in resource attributes do not take effect until you take the service group offline and bring it online. Therefore, you cannot access the Web Console using the new IP address, but you can from the previous address.

To reconfigure the virtual IP address:

1. Take the VCSweb and webip resources offline.
2. Change the address attribute of the webip resource.
3. Bring the VCSweb and webip resources online.

Flashing colors appear on Netscape while switching between Cluster Manager and other open windows.

Recommended Action: If there are flashes of color while viewing Cluster Manager on Netscape Navigator 4.7 or later, it is mostly likely a color-mapping issue. Set the display to 256 colors or a higher on the host machine where the GUI is being viewed to ensure best color and clarity.

"The object type specified is invalid. It should be one of cluster, group, type, resource, or system."

Recommended Action: This error (#W10002) occurs if the page URL points to a VCS object that does not exist or was deleted. If you typed the URL, verify the URL is correct. Names of VCS objects are case-sensitive: the object name in the URL must be entered in the correct case. If you clicked a link and got this error, refresh the page and retry. If you are still unsuccessful, contact Symantec Technical Support.

"The specified group does not exist or has been deleted."

Recommended Action: This error (#W10003) indicates the service group whose information you tried to access does not exist, or was deleted. If you typed the URL, verify the URL is correct. If you clicked a link to get information about the service group, verify the service group exists. Refresh the display to get current information.

"The specified system does not exist or has been deleted."

Recommended Action: This error (#W10004) indicates the system whose information you tried to access does not exist, or was deleted. If you typed the URL, verify the URL is correct. If you clicked a link to get information about the system, verify the system exists. Refresh the display to get current information.

"The specified resource type does not exist or has been deleted."

Recommended Action: This error (#W10005) indicates the resource type whose information you tried to access does not exist, or was deleted. If you typed the URL, verify the URL is correct. If you clicked a link to get information about the resource type, verify the resource type exists. Refresh the display to get current information.

"The specified resource does not exist or has been deleted."

Recommended Action: This error (#W10007) indicates the resource whose information you tried to access does not exist, or was deleted. If you typed the URL, verify the URL is correct. If you clicked a link to get information about the resource type, verify the resource exists. Refresh the display to get current information.

"Retrieving data from the VCS engine. Please try after some time."

Recommended Action: This error (#R10001) indicates a "snapshot" of the VCS engine, HAD, is being taken. Wait a few moments then retry the operation.

"Could not log on to the VCS engine."

Recommended Action: This error (#R10002) indicates Cluster Manger (Web Console) could not connect to the VCS engine. Wait a few moments then retry the operation.

"Cannot monitor VCS QuickStart."

Recommended Action: This error (R10005) indicates you tried to connect to a cluster configured by VCS QuickStart. Cluster Manager (Web Console) cannot connect to VCS QuickStart. Use the VCS QuickStart Web graphical user interface instead.

"The user could not be authenticated at this time. This could be because a snapshot of the VCS Server is being taken currently."

Recommended Action: This error (#H10001) indicates a snapshot of the VCS engine is being taken. Wait a few moments then retry the operation.

"The URL you specified can be accessed only if you are logged on."

Recommended Action: This error (#G10001) indicates you tried to access a page that requires authentication. Log on to VCS and retry the operation.

Troubleshooting VCS configuration backup and restore

This section cites the problem you may encounter when using the hasnap command to backup and restore VCS configuration files.

Error connecting to remote nodes in the cluster.

The hasnap command is a distributed command in the sense that it tries to backup and restore files from all cluster nodes in a single session. It needs to establish connection with all cluster nodes from the node where the command is executed. The connection may fail for one of the following reasons:

- The hasnap command retrieves the list of cluster nodes from the llthosts configuration file. However, the node names in this file may not always be DNS resolvable, in which case the command cannot establish connection with the remote nodes.

  Recommended Action:. For each node in the cluster, map the VCS node names to the actual DNS-resolvable names using the hasnap configuration file.

- The hasnap command uses the VCS Command Server Daemon running on the remote nodes to establish connection. The connection fails if the Daemon is not running on the remote node.

  Recommended Action: Verify the VCS Command Server Daemon is running on all cluster nodes.

- The remote node might be currently down or unreachable.

  Recommended Action: Run the hasnap command again after the bringing the remote node online.

Troubleshooting and recovery for global clusters

This section describes the concept of disaster declaration and provides troubleshooting tips for configurations using global clusters.

Disaster declaration

When a cluster in a global cluster transitions to the FAULTED state because it can no longer be contacted, failover executions depend on whether the cause was due to a split-brain, temporary outage, or a permanent disaster at the remote cluster.

If you choose to take action on the failure of a cluster in a global cluster, VCS prompts you to declare the type of failure.

- *Disaster*, implying permanent loss of the primary data center
- *Outage,* implying the primary may return to its current form in some time
- *Disconnect*, implying a split-brain condition; both clusters are up, but the link between them is broken
- *Replica*, implying that data on the takeover target has been made consistent from a backup source and that the RVGPrimary can initiate a takeover when the service group is brought online. This option applies to VVR environments only.

You can select the groups to be failed over to the local cluster, in which case VCS brings the selected groups online on a node based on the group's FailOverPolicy attribute. It also marks the groups as being offline in the other cluster. If you do not select any service groups to fail over, VCS takes no action except implicitly marking the service groups as offline on the downed cluster.

Lost heartbeats and the inquiry mechanism

The loss of internal and all external heartbeats between any two clusters indicates that the remote cluster is faulted, or that all communication links between the two clusters are broken (a wide-area split-brain).

A configuration with more than two clusters must distinguish between the two (Systems A and B) by querying the remaining clusters to confirm the remote cluster to which heartbeats have been lost is truly down. This mechanism is referred to as "Inquiry." If in a two-cluster configuration a connector loses all heartbeats (internal and external) to the other connector, it must consider the remote cluster faulted. If there are more than two clusters and a connector loses all heartbeats to a second cluster, it queries the remaining connectors regarding their "view" of the cluster in question before declaring it faulted. If the other connectors view the cluster as running (a negative inquiry), the querying connector transitions the cluster to the UNKNOWN state, a process that minimizes false cluster faults. If all connectors report that the cluster is faulted (a positive inquiry), the querying connector also considers it faulted and transitions the remote cluster state to FAULTED.

VCS alerts

VCS alerts are identified by the alert ID, which is comprised of the following elements:

- `alert_type`—The type of the alert, described in "Types of alerts."
- `cluster`—The cluster on which the alert was generated
- `system`—The system on which this alert was generated
- `object`—The name of the VCS object for which this alert was generated. This could be a cluster or a service group.

Alerts are generated in the following format:

```
alert_type-cluster-system-object
```

For example:

```
GNOFAILA-Cluster1-oracle_grp
```

This is an alert of type GNOFAILA generated on cluster Cluster1 for the service group oracle_grp.

Types of alerts

VCS generates the following types of alerts.

- CFAULT—Indicates that a cluster has faulted
- GNOFAILA—Indicates that a global group is unable to fail over within the cluster where it was online. This alert is displayed if the ClusterFailOverPolicy attribute is set to Manual and the wide-area connector (wac) is properly configured and running at the time of the fault.
- GNOFAIL—Indicates that a global group is unable to fail over to any system within the cluster or in a remote cluster.
  Some reasons why a global group may not be able to fail over to a remote cluster:
  - The ClusterFailOverPolicy is set to either Auto or Connected and VCS is unable to determine a valid remote cluster to which to automatically the group over.
  - The ClusterFailOverPolicy attribute is set to Connected and the cluster in which the group has faulted cannot communicate with one ore more remote clusters in the group's ClusterList.
  - The wide-area connector (wac) is not online or is incorrectly configured in the cluster in which the group has faulted

Managing alerts

Alerts require user intervention. You can respond to an alert in the following ways:

- If the reason for the alert can be ignored, use the Alerts dialog box in the Java or Web consoles or the `haalert` command to delete the alert. You must provide a comment as to why you are deleting the alert; VCS logs the comment to engine log.
- Take an action on administrative alerts that have actions associated with them. You can do so using the Java or Web consoles. See “Actions associated with alerts” for more information.
- VCS deletes or *negates* some alerts when a negating event for the alert occurs. See “Negating events” for more information.

An administrative alert will continue to live if none of the above actions are performed and the VCS engine (HAD) is running on at least one node in the cluster. If HAD is not running on any node in the cluster, the administrative alert is lost.

Actions associated with alerts

This section describes the actions you can perform from the Java and the Web consoles on the following types of alerts:

- CFAULT—When the alert is presented, clicking **Take Action** guides you through the process of failing over the global groups that were online in the cluster before the cluster faulted.
- GNOFAILA—When the alert is presented, clicking **Take Action** guides you through the process of failing over the global group to a remote cluster on which the group is configured to run.
- GNOFAIL—There are no associated actions provided by the consoles for this alert

Negating events

VCS deletes a CFAULT alert when the faulted cluster goes back to the running state

VCS deletes the GNOFAILA and GNOFAIL alerts in response to the following events:

- The faulted group's state changes from FAULTED to ONLINE.
- The group's fault is cleared.
- The group is deleted from the cluster where alert was generated.

VCS utilities

VCS provides several utilities that address common issues, however, *you must use them with extreme caution*. For best results, contact Symantec Technical Support prior to using the utilities described below.

The getcomms utility

The getcomms utility collects and saves information related to the private network. The information can be used by Technical Support to debug network-related issues.

To run the getcomms utility

Run getcomms using the Perl executables provided with VCS.

```
C:\> "VRTS_HOME\VRTSPerl\bin\perl.exe" getcomms.pl
        [-option]
```

The variable *VRTS_HOME* represents the Veritas installation directory, typically `C:\Program Files\VERITAS`. If you chose the default installation paths, use the following command to run getcomms:

```
C:\> "C:\Program Files\VERITAS\VRTSPerl\bin\perl.exe" getcomms.pl
        [-option]
```

getcomms options

Use the following options to limit the diagnostic information to specific components.

Table 23-3 getcomms options

Options	Action
-local	Retrieves and dumps information about the local system
-remote	Retrieves and dumps information about all live systems in the cluster
-stuck	Prints the message queue
-d *logdir*	Dumps information at the directory specified by the variable *logdir*

Log location

The getcomms utility writes the output to the directory *%temp%*\commslog.*timestamp* where *%temp%* is a system-defined environment variable and *timestamp* represents the time the log was taken.

The hagetcf utility

The hagetcf utility retrieves and writes detailed diagnostic information about the VCS configuration. The information can be used by Technical Support to debug configuration-related issues.

To access hagetcf, type:

```
C:\> hagetcf [-option]
```

Running hagetcf displays output similar to the example below:

```
    Veritas Cluster Server(TM) 5.0 for Windows 2000/2003 Diagnostics
    Copyright (C) 2007 Symantec Corporation. All rights reserved.
    Dumping output to: C:\Program Files\Veritas\Cluster
Server\hagetcf
```

Log location

By default, hagetcf writes the output to the directory *%VCS_HOME%*\hagetcf, where *%VCS_HOME%* is the VCS installation directory, typically C:\Program Files\VERITAS\Cluster Server\.

Options for the hagetcf utility

Use the following options with the hagetcf command to limit the diagnostic information to specific components.

Table 23-4 hagetcf command options

Options	Action
-app	Dumps the application event log.
-sec	Dumps the security event log.
-sys	Dumps the system event log.
-allevt	Dumps all event logs.
-conf	Dumps the VCS config directory.
-log	Dumps the VCS log directory.
-ldf	Dumps the VCS ldf directory.
-lock	Dumps the lock directory.
-triggers	Dump all files from the VCS triggers directory.
-alldir	Dumps the config, log, ldf, and lock directories.

Table 23-4 hagetcf command options (Continued)

Options	Action

-ha	Dumps the output of the following commands:
	`hares -display -all`
	`hagrp -display -all`
	`hasys -display`
	`getcomms.pl`

-state	Dumps the following system state information: ■ Dr. Watson log ■ Drive signature information from the havol utility ■ Network information, including NICs, ipconfig, and network-related registry entries ■ The VERITAS registry key ■ Output of the nbstat and the netstat commands ■ Licensing information ■ Disk and volume information ■ SCSI and Fibre configuration information ■ Server configuration information ■ Service and device state and configuration information ■ Processes running on the system ■ Information about the privileges of the current user ■ Information about products installed on the system
-haver	Dumps version information about all VCS binaries.
-nogetcomms	Excludes the output of the getcomms.pl command.
-sql	Dumps information about SQL Server and the VCS agent for SQL Server.
-exch	Dumps information about Exchange Server and the VCS agent for Exchange Server.
-iis	Dumps IIS information.
-allagents	Dumps information about all enterprise agents.
-vxlog	Dumps diagnostic information about other Veritas products.
-islog	Dumps installation log.
-o *<outdir>*	Dumps hagetcf output to *<outdir>*.
-? or -help	Displays the command's usage information.

Note: If you do not specify any options, the command retrieves diagnostic information with the following options: `-app, -sys, -ha, -log, -lock, -conf, -state, -islog, -trigger`

The NICTest utility

The NICTest utility determines whether a NIC maintains its connection status in a system-defined variable. The utility helps configuring NICs under VCS.

Using NICTest

To perform the NIC test

1. At the command prompt, type:

   ```
   C:\> NICTest adapter_macaddress
   ```

   The variable *adapter_macaddress* represents the physical address of the adapter. You can retrieve the MAC address using the `ipconfig -all` command. The utility displays an error message if the entered MAC address is invalid or if it cannot find the specified adapter.

   The utility prompts you to ensure that the NIC is connected:

   ```
   Please ensure that the NIC is connected to the network and
   press Enter.
   ```

2. Press Return. The utility prompts you to disconnect the NIC.
3. Disconnect the NIC and press Return. The system prompts you to connect the NIC.
4. Connect the NIC and press Return.

   If the NIC does not maintain its connection status, the following message is displayed:

   ```
   NIC <adapter_macaddress> does not maintain the connection
   status.
   ```

   If the NIC maintains its connection status, the following message is displayed:

   ```
   NIC <adapter_macaddress> maintains the connection status.
   Please set the UseConnectionStatus attribute for this
   resource to True.
   ```

The VCSRegUtil utility

If an application is run outside of VCS, registry changes are not logged to the shared disk. VCS provides a utility, VCSRegUtil.exe, that marks the system on which registry changes are made outside of the VCS environment.

If a system is marked by the VCSRegUtil utility, the agent detects registry changes when VCS is started. The agent then logs changes to the shared disk. Therefore, you must run the VCSRegUtil.exe utility whenever you run the clustered application outside of VCS. For example, you must use it when issuing the command `hastop -all -force` to take all resources offline and run the application outside the VCS environment. The utility also unmarks a previously marked system. When the resource is brought online on a system marked by this utility, the agent unmarks the system.

Note: If a system is marked using VCSRegUtil.exe, and if the attribute RestoreLocally is set to 1, the system marking overrides the RestoreLocally attribute and registry changes are not applied to the system when it is brought online.

To mark a system

At the command prompt, type:

```
C:\> VCSRegUtil /RESOURCE=RegRepResourceName /MARK
```

To unmark a system

At the command prompt, type:

```
C:\> VCSRegUtil /RESOURCE=RegRepResourceName /UNMARK
```

The havol utility

The havol utility provides the following options:

- -getdrive: Retrieves information about the disk and stores it in a file called DriveInfo.txt in the same path from where you executed the command.
- -scsitest: Reserves and releases disks. It helps troubleshoot issues related to SCSI reservation.

Syntax

```
C:\> havol -scsitest <options>
C:\> havol -getdrive [-details]
```

Using the -getdrive option

At the command prompt, type:

```
C:\> havol -getdrive
```

For detailed information about the disk, type:

```
C:\> havol -getdrive -details
```

The utility retrieves information about the disk and stores it in a file called DriveInfo.txt in the same path from where you executed the command.

Sample file contents include:

```
Detailed Information about Fixed Disks with valid volumes in the
system : VCSW2K112J

Harddisk Number  = 1
Harddisk Type    = Basic Disk
Disk Signature   = 130349684
Valid Partitions = 3
Access Test      = SUCCESS
Partition Number = 3
Partition Type   = IFS Partition
Bootable         = NO
Partition Size   = 400.06 MB
WINNT style Name = \device\Harddisk1\Partition3
Target Name      = \Device\HarddiskVolume6
Device Name      =
\\?\Volume{03074b0e-b4d7-11d6-b5a9-00d0b7471a1f}\
DriveLetter      = Unassigned
DrivePath001     = F:\f1\

Partition Number = 2
Partition Type   = IFS Partition
Bootable         = NO
Partition Size   = 400.06 MB
WINNT style Name = \device\Harddisk1\Partition2
Target Name      = \Device\HarddiskVolume5
```

```
Device Name       =
\\?\Volume{03074af7-b4d7-11d6-b5a9-00d0b7471a1f}\
DriveLetter       = Unassigned
DrivePath001      = F:\f2\

Partition Number = 1
Partition Type    = IFS Partition
Bootable          = NO
Partition Size    = 4.01 GB
WINNT style Name = \device\Harddisk1\Partition1
Target Name       = \Device\HarddiskVolume4
Device Name       =
\\?\Volume{e587ddc7-8cee-11d6-b401-00d0b7471a1f}\
DriveLetter       = F:
MountPoint001     = F:\f2\ ->
\\?\Volume{03074af7-b4d7-11d6-b5a9-00d0b7471a1f}\
MountPoint002     = F:\f1\ ->
\\?\Volume{03074b0e-b4d7-11d6-b5a9-00d0b7471a1f}\
```

Using the -scsitest option

At the command prompt, type:

```
C:\> havol -scsitest [/option]
```

The variable *option* represents additional parameters for the command, listed in “Options for the -scsitest command” on page 678.

Caution: Reserving or releasing shared disks may cause the configured service groups to fail over to another system.

Retrieving the disk number

Most scsitest options require the disk number. To list the disk numbers visible from the system, type the following command:

```
C:\> havol -scsitest /L
```

Verify the disk signature to ensure you choose the correct disk number. If the required disk number or signature is not listed, try rescanning the SCSI bus. Type:

```
C:\> havol -scsitest /RESCAN
```

Options for the -scsitest command

Use the following options with the scsitest command to limit the diagnostic information to specific components.

Table 23-5 -scscitest command: basic options

Basic Options	Action
-ADDR:1	Gets the SCSI address of disk number 1.
-LISTDISKS or **-L**	Lists all visible disks.
-REL:1	Releases disk number 1.
-RES:1	Reserves disk number 1.
-RESCAN	Rescans the SCSI bus.
-RESET:1	Resets the disk number 1 (in ioctl mode).
-SIG:1	Retrieves the signature of disk number 1.

Table 23-6 -scscitest command: advanced options

Advanced Options	Action
-DISKCOUNT	Returns the total number of disks reserved persistently.
-PREL:1	Persistently releases disk number 1.
-PRES:1	Persistently reserves disk number 1.
-REMOVEALL	Removes all disks from persistent reservation.
-RESETPBI:1,0	Resets the port number 1 and path 0 by ioctl mode.

Table 23-6 -scscitest command: advanced options (Continued)

Advanced Options	Action
-RESETPBS:1,0	Resets the port number 1 and path 0 by SRB mode.
-STARTDRV	Starts the DiskRes driver.
-STOPDRV	Stops the DiskRes driver.
-VER	Retrieves the DiskRes.sys version.

The vmgetdrive utility

Use the VMGetDrive utility to retrieve information about the cluster disk groups and configured volumes.

To retrieve information about the cluster disk groups using the VMGetDrive utility

1. At the command prompt, from the path %VCS_HOME%\bin, type:

   ```
   %VCS_HOME%\bin> vmgetdrive
   ```

   The system retrieves information about the volume and stores it in a file called VMDriveInfo.txt in the same path from where you executed the command.

2. Open the file VMDriveInfo.txt using a text editor, and get the required information. Sample file contents include:

   ```
   There are 1 Imported Cluster Disk Groups

   DiskGroup Name = VCS1
   No of disks in diskgroup 'VCS1'  = 2
       Harddisk2
       Harddisk3

   No of volumes in diskgroup 'VCS1' = 2
   Volume Name = Stripe1
       Drive Letter = NONE
       File System  = NTFS
       Mount Point  = NONE
       Volume Label =
       Volume Type  = STRIPED
   Volume Name = Volume1
       Drive Letter = NONE
       File System  = NTFS
       Mount Point  = NONE
       Volume Label =
   ```

   Volume Type = CONCATENATED

Configuring the VCS HAD helper service manually

Use the HadHelper command to configure the VCS HAD Helper service manually.

Command syntax

```
    HADHelper /Install /User:<user_name> [/Password:<password>]
HADHelper /Uninstall
HADHelper /Configure /User:<user_name> [/Password:<password>]
HADHelper /ShowConfig
```

- If you do not specify a password for the `/Install` and `/Configure` options, the command prompts for a password.
- Specify the <*user_name*> as domain\user or user. If you do not append the domain name, the command assumes the user belongs to the current domain.
- Assign the privilege **Add workstation to domain** on the domain controller to the user.

Command options

Note that the following command options are case-insensitive.

Table 23-7 HadHelper command options

Options	Action
`/Install`	Installs the HADHelper service, configures the user for the service, assigns the required local security privileges to the user, and adds the user to the local administrators group. If the service already exists, the option configures the user for the service.
`/Uninstall`	Uninstalls the service, removes the local security privileges for the configured user, and removes the user from local administrators group. **Note:** Stop the service before running the command to uninstall the service.
`/Configure`	Changes the user for the service, assigns the required local security privileges to the user, and adds the user to local administrators group. The option also revokes the local security privileges of the previous user and removes the user from local administrators group.

Table 23-7 HadHelper command options (Continued)

Options	Action
`/ShowConfig`	Displays the user name, user SID, and the local security privileges held by the user.

Section VII

Appendixes

- Appendix A, “VCS user privileges—administration matrices” on page 685
- Appendix B, “Cluster and system states” on page 693
- Appendix C, “VCS attributes” on page 701
- Appendix D, “Administering Symantec Web Server” on page 745
- Appendix E, “Configuring LLT over UDP” on page 729
- Appendix F, “Handling concurrency violation in Any-to-Any configurations” on page 737

Appendix A

VCS user privileges—administration matrices

- About administration matrices
- Administration matrices

About administration matrices

In general, users with Guest privileges can execute the following command options: -display, -state, and -value.

Users with privileges for Group Operator and Cluster Operator can execute the following options: -online, -offline, and -switch.

Users with Group Administrator and Cluster Administrator privileges can execute the following options -add, -delete, and -modify.

See "About the VCS user privilege model" on page 63.

Administration matrices

Review the matrices in the following section to determine which command options can be executed within a specific user role. Checkmarks denote the command and option can be executed. A dash indicates they cannot.

Agent Operations (haagent)

Table A-1 lists agent operations and required privileges.

Table A-1 User privileges for agent operations

Agent Operation	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Start agent	–	–	–	✔	✔
Stop agent	–	–	–	✔	✔
Display info	✔	✔	✔	✔	✔
List agents	✔	✔	✔	✔	✔

Attribute Operations (haattr)

Table A-2 lists attribute operations and required privileges.

Table A-2 User privileges for attribute operations

Attribute Operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add	–	–	–	–	✔
Change default value	–	–	–	–	✔
Delete	–	–	–	–	✔
Display	✔	✔	✔	✔	✔

Cluster Operations (haclus, haconf)

Table A-3 lists cluster operations and required privileges.

Table A-3 User privileges for cluster operations

Cluster Operations	Cluster Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Display	✔	✔	✔	✔	✔
Modify	–	–	–	–	✔
Add	–	–	–	–	✔
Delete	–	–	–	–	✔
Declare	–	–	–	✔	✔
View state or status	✔	✔	✔	✔	✔
Update license					✔
Make configuration read-write	–	–	✔	–	✔
Save configuration	–	–	✔	–	✔
Make configuration read-only	–	–	✔	–	✔

Service group operations (hagrp)

Table A-4 lists service group operations and required privileges.

Table A-4 User privileges for service group operations

Service Group Operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add and delete	–	–	–	–	✔
Link and unlink	–	–	–	–	✔
Clear	–	✔	✔	✔	✔
Bring online	–	✔	✔	✔	✔
Take offline	–	✔	✔	✔	✔
View state	✔	✔	✔	✔	✔
Switch	–	✔	✔	✔	✔
Freeze/unfreeze	–	✔	✔	✔	✔
Freeze/unfreeze persistent	–	–	✔	–	✔
Enable	–	–	✔	–	✔
Disable	–	–	✔	–	✔
Modify	–	–	✔	–	✔
Display	✔	✔	✔	✔	✔
View dependencies	✔	✔	✔	✔	✔
View resources	✔	✔	✔	✔	✔
List	✔	✔	✔	✔	✔
Enable resources	–	–	✔	–	✔
Disable resources	–	–	✔	–	✔
Flush	–	✔	✔	✔	✔
Autoenable	–	✔	✔	✔	✔
Ignore	–	✔	✔	✔	✔

Heartbeat operations (hahb)

Table A-5 lists heartbeat operations and required privileges.

Table A-5 User privileges for heartbeat operations

Heartbeat Operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add	–	–	–	–	✔
Delete	–	–	–	–	✔
Make local	–	–	–	–	✔
Make global	–	–	–	–	✔
Display	✔	✔	✔	✔	✔
View state	✔	✔	✔	✔	✔
List	✔	✔	✔	✔	✔

Log operations (halog)

Table A-6 lists log operations and required privileges.

Table A-6 User privileges for log operations

Log operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Enable debug tags	–	–	–	–	✔
Delete debug tags	–	–	–	–	✔
Add messages to log file	–	–	–	✔	✔
Display	✔	✔	✔	✔	✔
Display log file info	✔	✔	✔	✔	✔

Resource operations (hares)

Table A-7 lists resource operations and required privileges.

Table A-7 User privileges for resource operations

Resource operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add	–	–	✔	–	✔
Delete	–	–	✔	–	✔
Make local	–	–	✔	–	✔
Make global	–	–	✔	–	✔
Link and unlink	–	–	✔	–	✔
Clear	–	✔	✔	✔	✔
Bring online	–	✔	✔	✔	✔
Take offline	–	✔	✔	✔	✔
Modify	–	–	✔	–	✔
View state	✔	✔	✔	✔	✔
Display	✔	✔	✔	✔	✔
View dependencies	✔	✔	✔	✔	✔
List, Value	✔	✔	✔	✔	✔
Probe	–	✔	✔	✔	✔
Override	–	–	✔	–	✔
Remove overrides	–	–	✔	–	✔
Run an action	–	✔	✔	✔	✔
Refresh info	–	✔	✔	✔	✔
Flush info	–	✔	✔	✔	✔

System operations (hasys)

Table A-8 lists system operations and required privileges.

Table A-8 User privileges for system operations

System operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add	–	–	–	–	✔
Delete	–	–	–	–	✔
Freeze and unfreeze	–	–	–	✔	✔
Freeze and unfreeze persistent	–	–	–	–	✔
Freeze and evacuate	–	–	–	–	✔
Display	✔	✔	✔	✔	✔
Start forcibly	–	–	–	–	✔
Modify	–	–	–	–	✔
View state	✔	✔	✔	✔	✔
List	✔	✔	✔	✔	✔
Update license	–	–	–	–	✔

Resource type operations (hatype)

Table A-9 lists resource type operations and required privileges.

Table A-9 User privileges for resource type operations

Resource type operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add	–	–	–	–	✔
Delete	–	–	–	–	✔
Display	✔	✔	✔	✔	✔
View resources	✔	✔	✔	✔	✔
Modify	–	–	–	–	✔
List	✔	✔	✔	✔	✔

User operations (hauser)

Table A-10 lists user operations and required privileges.

Table A-10 User privileges for user operations

User operations	Guest	Group Operator	Group Admin.	Cluster Operator	Cluster Admin.
Add	–	–	–	–	✔
Delete	–	–	–	–	✔
Update	–	✔	✔	✔	✔
Display	✔	✔	✔	✔	✔
List	✔	✔	✔	✔	✔
Modify privileges	–	–	✔	–	✔

Appendix B

Cluster and system states

- Remote cluster states
- System states

Remote cluster states

In global clusters, the “health” of the remote clusters is monitored and maintained by the wide-area connector process. The connector process uses heartbeats, such as ICMP, to monitor the state of remote clusters. The state is then communicated to HAD, which then uses the information to take appropriate action when required. For example, when a cluster is shut down gracefully, the connector transitions its local cluster state to EXITING and notifies the remote clusters of the new state. When the cluster exits and the remote connectors lose their TCP/IP connection to it, each remote connector transitions their view of the cluster to EXITED.

To enable wide-area network heartbeats, the wide-area connector process must be up and running. For wide-area connectors to connect to remote clusters, at least one heartbeat to the specified cluster must report the state as ALIVE.

There are three hearbeat states for remote clusters: HBUNKNOWN, HBALIVE, and HBDEAD.

The following table provides a list of VCS remote cluster states and their descriptions.

See "Examples of system state transitions" on page 699.

Table B-11 VCS state definitions

State	Definition
INIT	The initial state of the cluster. This is the default state.
BUILD	The local cluster is receiving the initial snapshot from the remote cluster.
RUNNING	Indicates the remote cluster is running and connected to the local cluster.
LOST_HB	The connector process on the local cluster is not receiving heartbeats from the remote cluster
LOST_CONN	The connector process on the local cluster has lost the TCP/IP connection to the remote cluster.
UNKNOWN	The connector process on the local cluster determines the remote cluster is down, but another remote cluster sends a response indicating otherwise.
FAULTED	The remote cluster is down.
EXITING	The remote cluster is exiting gracefully.
EXITED	The remote cluster exited gracefully.
INQUIRY	The connector process on the local cluster is querying other clusters on which heartbeats were lost.
TRANSITIONING	The connector process on the remote cluster is failing over to another node in the cluster.

Examples of cluster state transitions

- If a remote cluster joins the global cluster configuration, the other clusters in the configuration transition their “view” of the remote cluster to the RUNNING state:

  INIT -> BUILD -> RUNNING

- If a cluster loses all heartbeats to a remote cluster in the RUNNING state, inquiries are sent. If all inquiry responses indicate the remote cluster is actually down, the cluster transitions the remote cluster state to FAULTED:

  RUNNING -> LOST_HB -> INQUIRY -> FAULTED

- If at least one response does not indicate the cluster is down, the cluster transitions the remote cluster state to UNKNOWN:

  RUNNING -> LOST_HB -> INQUIRY -> UNKNOWN

- When the ClusterService service group, which maintains the connector process as highly available, fails over to another system in the cluster, the remote clusters transition their view of that cluster to TRANSITIONING, then back to RUNNING after the failover is successful:

  RUNNING -> TRANSITIONING -> BUILD -> RUNNING

- When a remote cluster in a RUNNING state is stopped (by taking the ClusterService service group offline), the remote cluster transitions to EXITED:

  RUNNING -> EXITING -> EXITED

System states

Whenever the VCS engine is running on a system, it is in one of the states described in the table below. States indicate a system’s current mode of operation. When the engine is started on a new system, it identifies the other systems available in the cluster and their states of operation. If a cluster system is in the state of RUNNING, the new system retrieves the configuration information from that system. Changes made to the configuration while it is being retrieved are applied to the new system before it enters the RUNNING state.

If no other systems are up and in the state of RUNNING or ADMIN_WAIT, and the new system has a configuration that is not invalid, the engine transitions to the state LOCAL_BUILD, and builds the configuration from disk. If the configuration is invalid, the system transitions to the state of STALE_ADMIN_WAIT.

The following table provides a list of VCS system states and their descriptions.

See "Examples of system state transitions" on page 699.

Table B-12 VCS system states

State	Definition
ADMIN_WAIT	The running configuration was lost. A system transitions into this state for the following reasons: ■ The last system in the RUNNING configuration leaves the cluster before another system takes a snapshot of its configuration and transitions to the RUNNING state. ■ A system in LOCAL_BUILD state tries to build the configuration from disk and receives an unexpected error from **hacf** indicating the configuration is invalid.
CURRENT_DISCOVER_WAIT	The system has joined the cluster and its configuration file is valid. The system is waiting for information from other systems before it determines how to transition to another state.
CURRENT_PEER_WAIT	The system has a valid configuration file and another system is doing a build from disk (LOCAL_BUILD). When its peer finishes the build, this system transitions to the state REMOTE_BUILD.
EXITING	The system is leaving the cluster.
EXITED	The system has left the cluster.
EXITING_FORCIBLY	An `hastop -force` command has forced the system to leave the cluster.
FAULTED	The system has left the cluster unexpectedly.
INITING	The system has joined the cluster. This is the initial state for all systems.
LEAVING	The system is leaving the cluster gracefully. When the agents have been stopped, and when the current configuration is written to disk, the system transitions to EXITING.
LOCAL_BUILD	The system is building the running configuration from the disk configuration.
REMOTE_BUILD	The system is building a running configuration that it obtained from a peer in a RUNNING state.

Table B-12 VCS system states

State	Definition
RUNNING	The system is an active member of the cluster.
STALE_ADMIN_WAIT	The system has an invalid configuration and there is no other system in the state of RUNNING from which to retrieve a configuration. If a system with a valid configuration is started, that system enters the LOCAL_BUILD state. Systems in STALE_ADMIN_WAIT transition to STALE_PEER_WAIT.
STALE_DISCOVER_WAIT	The system has joined the cluster with an invalid configuration file. It is waiting for information from any of its peers before determining how to transition to another state.
STALE_PEER_WAIT	The system has an invalid configuration file and another system is doing a build from disk (LOCAL_BUILD). When its peer finishes the build, this system transitions to the state REMOTE_BUILD.
UNKNOWN	The system has not joined the cluster because it does not have a system entry in the configuration.

Examples of system state transitions

- If VCS is started on a system, and if that system is the only one in the cluster with a valid configuration, the system transitions to the RUNNING state:

  INITING -> CURRENT_DISCOVER_WAIT -> LOCAL_BUILD -> RUNNING

- If VCS is started on a system with a valid configuration file, and if at least one other system is already in the RUNNING state, the new system transitions to the RUNNING state:

  INITING -> CURRENT_DISCOVER_WAIT -> REMOTE_BUILD -> RUNNING

- If VCS is started on a system with an invalid configuration file, and if at least one other system is already in the RUNNING state, the new system transitions to the RUNNING state:

  INITING -> STALE_DISCOVER_WAIT -> REMOTE_BUILD -> RUNNING

- If VCS is started on a system with an invalid configuration file, and if all other systems are in STALE_ADMIN_WAIT state, the system transitions to the STALE_ADMIN_WAIT state as shown below. A system stays in this state until another system with a valid configuration file is started.

  INITING -> STALE_DISCOVER_WAIT -> STALE_ADMIN_WAIT

- If VCS is started on a system with a valid configuration file, and if other systems are in the ADMIN_WAIT state, the new system transitions to the ADMIN_WAIT state.

  INITING -> CURRENT_DISCOVER_WAIT -> ADMIN_WAIT

- If VCS is started on a system with an invalid configuration file, and if other systems are in the ADMIN_WAIT state, the new system transitions to the ADMIN_WAIT state.

  INITING -> STALE_DISCOVER_WAIT -> ADMIN_WAIT

- When a system in RUNNING state is stopped with the `hastop` command, it transitions to the EXITED state as shown below. During the LEAVING state, any online system resources are taken offline. When all of the system's resources are taken offline and the agents are stopped, the system transitions to the EXITING state, then EXITED.

  RUNNING -> LEAVING -> EXITING -> EXITED

Appendix C

VCS attributes

- About attributes and their definitions
- Resource attributes
- Resource type attributes
- Service group attributes
- System attributes
- Cluster attributes
- Heartbeat attributes (for global clusters)

About attributes and their definitions

In addition to the attributes listed in this appendix, see the *Veritas Cluster Server Agent Developer's Guide*.

- You can modify the values of attributes labelled user-defined from the command line or graphical user interface, or by manually modifying the main.cf configuration file. You can change the default values to better suit your environment and enhance performance.
  When changing the values of attributes, be aware that VCS attributes interact with each other. After changing the value of an attribute, observe the cluster systems to confirm that unexpected behavior does not impair performance.
- The values of attributes labelled system use only are set by VCS and are read-only. They contain important information about the state of the cluster.
- The values labeled agent-defined are set by the corresponding agent and are also read-only.

See "About VCS attributes" on page 52.

Resource attributes

Table C-1 lists resource attributes.

Table C-1 Resource attributes

Resource Attributes	Description
ArgListValues (agent-defined)	List of arguments passed to the resource's agent on each system.This attribute is resource- and system-specific, meaning that the list of values passed to the agent depend on which system and resource they are intended. ■ Type and dimension: string-vector ■ Default: non-applicable.
AutoStart (user-defined)	Indicates the resource is brought online when the service group is brought online. ■ Type and dimension: boolean-scalar ■ Default: 1

Table C-1 Resource attributes

Resource Attributes	Description
ComputeStats (user-defined)	Indicates to agent framework whether or not to calculate the resource's monitor statistics. ■ Type and dimension: boolean-scalar ■ Default: 0
ConfidenceLevel (agent-defined)	Indicates the level of confidence in an online resource. Values range from 0–100. Note that some VCS agents may not take advantage of this attribute and may always set it to 0. Set the level to 100 if the attribute is not used. ■ Type and dimension: integer-scalar ■ Default: 0
Critical (user-defined)	Indicates whether a fault of this resource should trigger a failover of the entire group or not. If Critical is 0, the resource fault will not cause group failover. ■ Type and dimension: boolean-scalar ■ Default: 1
Enabled (user-defined)	Indicates agents monitor the resource. If a resource is created dynamically while VCS is running, you must enable the resource before VCS monitors it. For more information on how to add or enable resources, see the chapters on administering VCS from the command line and graphical user interfaces. When Enabled is set to 0, it implies a disabled resource. ■ Type and dimension: boolean-scalar ■ Default: If you specify the resource in main.cf prior to starting VCS, the default value for this attribute is 1, otherwise it is 0.

Table C-1 Resource attributes

Resource Attributes	Description
Flags (system use only)	Provides additional information for the state of a resource. Primarily this attribute raises flags pertaining to the resource. Values: NORMAL—Standard working order. RESTARTING —The resource faulted and that the agent is attempting to restart the resource on the same system. STATE UNKNOWN—The latest monitor call by the agent could not determine if the resource was online or offline. MONITOR TIMEDOUT —The latest monitor call by the agent was terminated because it exceeded the maximum time specified by the static attribute MonitorTimeout. UNABLE TO OFFLINE—The agent attempted to offline the resource but the resource did not go offline. This flag is also set when a resource faults and the clean entry point completes successfully, but the subsequent monitor hangs or is unable to determine resource status. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
Group (system use only)	String name of the service group to which the resource belongs. ■ Type and dimension: string-scalar ■ Default: Not applicable.

Table C-1 Resource attributes

Resource Attributes	Description
IState (system use only)	The internal state of a resource. In addition to the State attribute, this attribute shows to which state the resource is transitioning. Values: NOT WAITING—Resource is not in transition. WAITING TO GO ONLINE—Agent notified to bring the resource online but procedure not yet complete. WAITING FOR CHILDREN ONLINE—Resource to be brought online, but resource depends on at least one offline resource. Resource transitions to WAITING TO GO ONLINE when all children are online. WAITING TO GO OFFLINE—Agent notified to take the resource offline but procedure not yet complete. WAITING TO GO OFFLINE (propagate)—Same as above, but when completed the resource's children will also be offline. WAITING TO GO ONLINE (reverse)—Resource waiting to be brought online, but when it is online it attempts to go offline. Typically this is the result of issuing an offline command while resource was waiting to go online. WAITING TO GO OFFLINE (reverse/propagate)—Same as above, but resource propagates the offline operation. ■ Type and dimension: integer-scalar ■ Default: 1 NOT WAITING
LastOnline (system use only)	Indicates the system name on which the resource was last online. This attribute is set by VCS. ■ Type and dimension: string-scalar ■ Default: Not applicable
MonitorOnly (system use only)	Indicates if the resource can be brought online or taken offline. If set to 0, resource can be brought online or taken offline. If set to 1, resource can only be monitored. **Note:** This attribute can only be modified by the command `hagrp -freeze`. ■ Type and dimension: boolean-scalar ■ Default: 0

Table C-1 Resource attributes

Resource Attributes	Description
MonitorTimeStats (system use only)	Valid keys are Average and TS. Average is the average time taken by the monitor entry point over the last Frequency number of monitor cycles. TS is the timestamp indicating when the engine updated the resource's Average value. ■ Type and dimension: string-association ■ Default: Average = 0 TS = ""
Name (system use only)	Contains the actual name of the resource. ■ Type and dimension: string-scalar ■ Default: Not applicable.
Path (system use only)	Set to 1 to identify a resource as a member of a path in the dependency tree to be taken offline on a specific system after a resource faults. ■ Type and dimension: boolean-scalar ■ Default: 0
Probed (system use only)	Indicates whether the state of the resource has been determined by the agent by running the monitor entry point. ■ Type and dimension: boolean-scalar ■ Default: 0
ResourceInfo (system use only)	This attribute has three predefined keys: State: values are Valid, Invalid, or Stale Msg: output of the info entry point captured on stdout by the agent framework TS: timestamp indicating when the ResourceInfo attribute was updated by the agent framework ■ Type and dimension: string-association ■ Default: State = Valid Msg = "" TS = ""

Table C-1 Resource attributes

Resource Attributes	Description
ResourceOwner (user-defined)	Used for VCS email notification and logging. VCS sends email notification to the person designated in this attribute when an event occurs related to the resource. VCS also logs the owner name when an event occurs. ■ Type and dimension: string-scalar ■ Default: If ResourceOwner is not specified in main.cf, the default value is "unknown".
Signaled (system use only)	Indicates whether a resource has been traversed. Used when bringing a service group online or taking it offline. ■ Type and dimension: integer-association ■ Default: Not applicable.
Start (system use only)	Indicates whether a resource was started (the process of bringing it online was initiated) on a system. ■ Type and dimension: integer -scalar ■ Default: 0
State (system use only)	Resource state displays the state of the resource and the flags associated with the resource. (Flags are also captured by the Flags attribute.) This attribute and Flags present a comprehensive view of the resource's current state. Values: ONLINE OFFLINE FAULTED ONLINE|STATE UNKNOWN ONLINE|MONITOR TIMEDOUT ONLINE|UNABLE TO OFFLINE OFFLINE|STATE UNKNOWN FAULTED|RESTARTING A FAULTED resource is physically offline, though unintentionally. ■ Type and dimension: integer -scalar ■ Default: 0
TriggerEvent (system use only)	A flag that turns Events on or off. ■ Type and dimension: boolean-scalar ■ Default: 0

Resource type attributes

Table C-2 lists the resource type attributes.

You can override some static attributes for resource types.

See "Overriding resource type static attributes" on page 209.

For more information on any attribute listed below, see the chapter on setting agent parameters in the *Veritas Cluster Server Agent Developer's Guide*.

Table C-2 Resource type attributes

Resource type attributes	Description
ActionTimeout (user-defined)	Timeout value for the Action entry point. ■ Type and dimension: integer-scalar ■ Default: 30 seconds
AgentClass (user-defined)	Indicates the scheduling class for the VCS agent process. ■ Type and dimension: string-scalar ■ Default: TS
AgentDirectory (user-defined)	Complete path of the directory in which the agent binary and scripts are located. Use this attribute in conjunction with the AgentFile attribute to specify a different location or different binary for the agent. ■ Type and dimension: string-scalar ■ Default = ""
AgentFailedOn (system use only)	A list of systems on which the agent for the resource type has failed. ■ Type and dimension: string-keylist ■ Default: Not applicable.
AgentFile (user-defined)	Complete name and path of the binary for an agent. If you do not specify a value for this attribute, VCS uses the agent binary at the path defined by the AgentDirectory attribute. ■ Type and dimension: string-scalar ■ Default = ""
AgentPriority (user-defined)	Indicates the priority in which the agent process runs. ■ Type and dimension: string-scalar ■ Default: 0

Table C-2 Resource type attributes

Resource type attributes	Description
AgentReplyTimeout (user-defined)	The number of seconds the engine waits to receive a heartbeat from the agent before restarting the agent. ■ Type and dimension: integer-scalar ■ Default: 130 seconds
AgentStartTimeout (user-defined)	The number of seconds after starting the agent that the engine waits for the initial agent "handshake" before restarting the agent. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
ArgList (user-defined)	An ordered list of attributes whose values are passed to the open, close, online, offline, monitor, clean, info, and action entry points. ■ Type and dimension: string-vector ■ Default: Not applicable.
AttrChangedTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the attr_changed entry point must complete or be terminated. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
CleanTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the clean entry point must complete or else be terminated. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
CloseTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the close entry point must complete or else be terminated. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
ConfInterval (user-defined) **Note:** This attribute can be overridden.	When a resource has remained online for the specified time (in seconds), previous faults and restart attempts are ignored by the agent. (See ToleranceLimit and RestartLimit attributes for details.) ■ Type and dimension: integer-scalar ■ Default: 600 seconds

Table C-2 Resource type attributes

Resource type attributes	Description
FaultOnMonitorTimeouts (user-defined) **Note:** This attribute can be overridden.	When a monitor times out as many times as the value specified, the corresponding resource is brought down by calling the clean entry point. The resource is then marked FAULTED, or it is restarted, depending on the value set in the RestartLimit attribute. When FaultOnMonitorTimeouts is set to 0, monitor failures are not considered indicative of a resource fault. A low value may lead to spurious resource faults, especially on heavily loaded systems. ■ Type and dimension: integer-scalar ■ Default: 4
FireDrill (user-defined)	Specifies whether or not fire drill is enabled for resource type. If set to 1, fire drill is enabled. If set to 0, it is disabled. ■ Type and dimension: boolean-scalar ■ Default: 0
InfoInterval (user-defined)	Duration (in seconds) after which the info entry point is invoked by the agent framework for ONLINE resources of the particular resource type. If set to 0, the agent framework does not periodically invoke the info entry point. To manually invoke the info entry point, use the command `hares -refreshinfo`. If the value you designate is 30, for example, the entry point is invoked every 30 seconds for all ONLINE resources of the particular resource type. ■ Type and dimension: integer-scalar ■ Default: 0
InfoTimeout (user-defined)	Timeout value for info entry point. If entry point does not complete by the designated time, the agent framework cancels the entry point's thread. ■ Type and dimension: integer-scalar ■ Default: 30 seconds

Table C-2 Resource type attributes

Resource type attributes	Description
LogDbg (user-defined)	Indicates the debug severities enabled for the resource type or agent framework. Debug severities used by the agent entry points are in the range of DBG_1–DBG_21. The debug messages from the agent framework are logged with the severities DBG_AGINFO, DBG_AGDEBUG and DBG_AGTRACE, representing the least to most verbose. ■ Type and dimension: string-keylist ■ Default: {} (none)
LogFileSize (user-defined)	Specifies the size (in bytes) of the agent log file. Minimum value is 65536 bytes. Maximum value is 134217728 bytes (128MB). ■ Type and dimension: integer-scalar ■ Default: 33554432 (32MB)
MonitorInterval (user-defined) **Note:** This attribute can be overridden.	Duration (in seconds) between two consecutive monitor calls for an ONLINE or transitioning resource. A low value may impact performance if many resources of the same type exist. A high value may delay detection of a faulted resource. ■ Type and dimension: integer-scalar ■ Default: 60 seconds

Table C-2 Resource type attributes

Resource type attributes	Description
MonitorStatsParam (user-defined)	Stores the required parameter values for calculating monitor time statistics. `static str MonitorStatsParam = {Frequency = 10, ExpectedValue = 3000, ValueThreshold = 100, AvgThreshold = 40}` Frequency: The number of monitor cycles after which the average monitor cycle time should be computed and sent to the engine. If configured, the value for this attribute must be between 1 and 30. The value 0 indicates that the monitor cycle ti me should not be computed. Default=0. ExpectedValue: The expected monitor time in milliseconds for all resources of this type. Default=100. ValueThreshold: The acceptable percentage difference between the expected monitor cycle time (ExpectedValue) and the actual monitor cycle time. Default=100. AvgThreshold: The acceptable percentage difference between the benchmark average and the moving average of monitor cycle times. Default=40. ■ Type and dimension: integer-association ■ Default: Different value for each parameter.
MonitorTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the monitor entry point must complete or else be terminated. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
NumThreads (user-defined)	Number of threads used within the agent process for managing resources. This number does not include threads used for other internal purposes. If the number of resources being managed by the agent is less than or equal to the NumThreads value, only that many number of threads are created in the agent. Addition of more resources does not create more service threads. Similarly deletion of resources causes service threads to exit. Thus, setting NumThreads to 1 forces the agent to just use 1 service thread no matter what the resource count is. The agent framework limits the value of this attribute to 30. ■ Type and dimension: integer-scalar ■ Default: 10

Table C-2 Resource type attributes

Resource type attributes	Description
OfflineMonitorInterval (user-defined) **Note:** This attribute can be overridden.	Duration (in seconds) between two consecutive monitor calls for an OFFLINE resource. If set to 0, OFFLINE resources are not monitored. ■ Type and dimension: integer-scalar ■ Default: 300 seconds
OfflineTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the offline entry point must complete or else be terminated. ■ Type and dimension: integer-scalar ■ Default: 300 seconds
OnlineRetryLimit (user-defined) **Note:** This attribute can be overridden.	Number of times to retry the `online` operation if the attempt to online a resource is unsuccessful. This parameter is meaningful only if the clean operation is implemented. ■ Type and dimension: integer-scalar ■ Default: 0
OnlineTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the online entry point must complete or else be terminated. ■ Type and dimension: integer-scalar ■ Default: 300 seconds
OnlineWaitLimit (user-defined) **Note:** This attribute can be overridden.	Number of monitor intervals to wait for the resource to come online after completing the online procedure. Increase the value of this attribute if the resource is likely to take a longer time to come online. ■ Type and dimension: integer-scalar ■ Default: 2
OpenTimeout (user-defined) **Note:** This attribute can be overridden.	Maximum time (in seconds) within which the open entry point must complete or else be terminated. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
Operations (user-defined)	Indicates valid operations for resources of the resource type. Values are OnOnly (can online only), OnOff (can online and offline), None (cannot online or offline). ■ Type and dimension: string-scalar ■ Default: OnOff

Table C-2 Resource type attributes

Resource type attributes	Description
RestartLimit (user-defined) **Note:** This attribute can be overridden.	Number of times to retry bringing a resource online when it is taken offline unexpectedly and before VCS declares it FAULTED. ■ Type and dimension: integer-scalar ■ Default: 0
ScriptClass (user-defined)	Indicates the scheduling class of the script processes (for example, online) created by the agent. ■ Type and dimension: string-scalar ■ Default: TS
ScriptPriority (user-defined)	Indicates the priority of the script processes created by the agent. ■ Type and dimension: string-scalar ■ Default: 0
SourceFile (user-defined)	File from which the configuration is read. Always set to.\types.cf. Make sure the path exists on all nodes before configuring this attribute. ■ Type and dimension: string-scalar ■ Default: .\types.cf
SupportedActions (user-defined)	Valid action tokens for the resource type. ■ Type and dimension: string-vector ■ Default: {}
ToleranceLimit (user-defined) **Note:** This attribute can be overridden.	After a resource goes online, the number of times the monitor entry point should return OFFLINE before declaring the resource FAULTED. A large value could delay detection of a genuinely faulted resource. ■ Type and dimension: integer-scalar ■ Default: 0

Service group attributes

Table C-3 lists the service group attributes.

Table C-3 Service group attributes

Service Group Attributes	Definition
ActiveCount (system use only)	Number of resources in a service group that are active (online or waiting to go online). When the number drops to zero, the service group is considered offline. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
AdministratorGroups (user-defined)	List of operating system user account groups that have administrative privileges on the service group. This attribute is valid only in secure clusters. ■ Type and dimension: string-keylist ■ Default: ""
Administrators (user-defined)	List of VCS users with privileges to administer the group. **Note:** A Group Administrator can perform all operations related to a specific service group, but cannot perform generic cluster operations. See "About the VCS user privilege model" on page 63. ■ Type and dimension: string-keylist ■ Default: ""
Authority (user-defined)	Indicates whether or not the local cluster is allowed to bring the service group online. If set to 0, it is not, if set to 1, it is. Only one cluster can have this attribute set to 1 for a specific global group. See "Serialization–The authority attribute" on page 547. ■ Type and dimension: integer-scalar ■ Default: 0
AutoDisabled (system use only)	Indicates that VCS does not know the status of a service group (or specified system for parallel service groups). This could occur because the group is not probed (on specified system for parallel groups) in the SystemList attribute. Or the VCS engine is not running on a node designated in the SystemList attribute, but the node is visible. ■ Type and dimension: boolean-scalar ■ Default: 0

Table C-3 Service group attributes

Service Group Attributes	Definition
AutoFailOver (user-defined)	Indicates whether VCS initiates an automatic failover if the service group faults. ■ Type and dimension: boolean-scalar ■ Default: 1 (enabled)
AutoRestart (user-defined)	Restarts a service group after a faulted persistent resource becomes online. See "About service group dependencies" on page 492. **Note:** This attribute applies only to service groups containing persistent resources. ■ Type and dimension: boolean-scalar ■ Default: 1 (enabled)
AutoStart (user-defined)	Designates whether a service group is automatically started when VCS is started. ■ Type and dimension: boolean-scalar ■ Default: 1 (enabled)
AutoStartIfPartial (user-defined)	Indicates whether to initiate bringing a service group online if the group is probed and discovered to be in a PARTIAL state when VCS is started. ■ Type and dimension: boolean-scalar ■ Default: 1 (enabled)
AutoStartList (user-defined)	List of systems on which, under specific conditions, the service group will be started with VCS (usually at system boot). For example, if a system is a member of a failover service group's AutoStartList attribute, and if the service group is not already running on another system in the cluster, the group is brought online when the system is started. VCS uses the AutoStartPolicy attribute to determine the system on which to bring the service group online. **Note:** For the service group to start, AutoStart must be enabled and Frozen must be 0. Also, beginning with 1.3.0, you must define the SystemList attribute prior to setting this attribute. ■ Type and dimension: string-keylist ■ Default: "" (none)

Table C-3 Service group attributes

Service Group Attributes	Definition
AutoStartPolicy (user-defined)	Sets the policy VCS uses to determine on which system to bring a service group online if multiple systems are available. This attribute has three options: Order (default)—Systems are chosen in the order in which they are defined in the AutoStartList attribute. Load—Systems are chosen in the order of their capacity, as designated in the AvailableCapacity system attribute. System with the highest capacity is chosen first. Priority—Systems are chosen in the order of their priority in the SystemList attribute. Systems with the lowest priority is chosen first. ■ Type and dimension: string-scalar ■ Default: Order
ClusterFailOverPolicy (user-defined)	Determines how a global service group behaves when a cluster faults. The attribute can take the following values: Manual—The group does not fail over to another cluster automatically. Auto—The group fails over to another cluster automatically if it is unable to fail over within the local cluster, or if the entire cluster faults. Connected—The group fails over automatically to another cluster only if it is unable to fail over within the local cluster. ■ Type and dimension: string-scalar ■ Default: Manual
ClusterList (user-defined)	Specifies the list of clusters on which the service group is configured to run. ■ Type and dimension: integer-association ■ Default: Not applicable.
CurrentCount (system use only)	Number of systems on which the service group is active. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
DeferAutoStart (system use only)	Indicates whether HAD defers the auto-start of a local group in case the global cluster is not fully connected. ■ Type and dimension: boolean-scalar ■ Default: Not applicable

Table C-3 Service group attributes

Service Group Attributes	Definition
Enabled (user-defined)	Indicates if a group can be failed over or brought online. If any of the local values are disabled, the group is disabled. ■ Type and dimension: boolean-scalar ■ Default: 1 (enabled)
Evacuate (user-defined)	Indicates if VCS initiates an automatic failover when user issues `hastop -local -evacuate`. ■ Type and dimension: boolean-scalar ■ Default: 1
Evacuating (system use only)	Indicates the node ID from which the service group is being evacuated. ■ Type and dimension: integer-scalar ■ Default: Not applicable
Failover (system use only)	Indicates service group is in the process of failing over. ■ Type and dimension: boolean-scalar ■ Default: Not applicable
FailOverPolicy (user-defined)	Sets the policy VCS uses to determine which system a group fails over to if multiple systems exist. This attribute can take the following values: Priority—The system defined as the lowest priority in the SystemList attribute is chosen. Load—The system defined with the least value in the system's Load attribute is chosen. RoundRobin—Systems are chosen according to how many active service groups they are hosting. The system with the least number of active service groups is chosen first. ■ Type and dimension: string-scalar ■ Default: Priority

Table C-3 Service group attributes

Service Group Attributes	Definition
FaultPropagation (user-defined)	Specifies if VCS should propagate the fault up to parent resources and take the entire service group offline when a resource faults. The value 1 indicates that when a resource faults, VCS fails over the service group, if the group's AutoFailOver attribute is set to 1. If The value 0 indicates that when a resource faults, VCS does not take other resources offline, regardless of the value of the Critical attribute. The service group does not fail over on resource fault. ■ Type and dimension: boolean-scalar ■ Default: 1
FromQ (system use only)	Indicates the system name from which the service group is failing over. This attribute is specified when service group failover is a direct consequence of the group event, such as a resource fault within the group or a group switch. ■ Type and dimension: string-association ■ Default: Not applicable
Frozen (user-defined)	Disables all actions, including autostart, online and offline, and failover, except for monitor actions performed by agents. (This convention is observed by all agents supplied with VCS.) ■ Type and dimension: boolean-scalar ■ Default: 0 (not frozen)
GroupOwner (user-defined)	VCS sends email notification to the person designated in this attribute when an event occurs related to the service group. VCS also logs the owner name when an event occurs. ■ Type and dimension: string-scalar ■ Default: If GroupOwner is not specified in main.cf, the default value is "unknown".
Guests (user-defined)	List of operating system user account groups that have Guest privileges on the service group. This attribute is valid only in secure clusters. ■ Type and dimension: string-keylist ■ Default: ""

Table C-3 Service group attributes

Service Group Attributes	Definition
IntentOnline (system use only)	Indicates whether to keep service groups online or offline. VCS sets this attribute to 1 if an attempt has been made to bring the service group online. For failover groups, VCS sets this attribute to 0 when the group is taken offline. For parallel groups, it is set to 0 for the system when the group is taken offline or when the group faults and can fail over to another system. VCS sets this attribute to 2 for failover groups if VCS attempts to autostart a service group; for example, attempting to bring a service group online on a system from AutoStartList. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
LastSuccess (system use only)	Indicates the time when service group was last brought online. ■ Type and dimension: integer-scalar ■ Default: Not applicable
Load (user-defined)	Integer value expressing total system load this group will put on a system. For example, the administrator may assign a value of 100 to a large production SQL and 15 to a Web server. ■ Type and dimension: integer-scalar ■ Default: 0
ManageFaults (user-defined)	Specifies if VCS manages resource failures within the service group by calling the Clean entry point for the resources. This attribute can take the following values. NONE—VCS does not call the Clean entry point for any resource in the group. User intervention is required to handle resource faults. See"Controlling Clean behavior on resource faults" on page 454. ■ Type and dimension: string-scalar ■ Default: ALL
ManualOps (user-defined)	Indicates if manual operations are allowed on the service group. ■ Type and dimension: boolean-scalar ■ Default = 1 (enabled)

Table C-3 Service group attributes

Service Group Attributes	Definition
MigrateQ (system use only)	Indicates the system from which the service group is migrating. This attribute is specified when group failover is an indirect consequence (in situations such as a system shutdown or another group faults and is linked to this group). ■ Type and dimension: string-association ■ Default: Not applicable
NumRetries (system use only)	Indicates the number of attempts made to bring a service group online. This attribute is used only if the attribute OnlineRetryLimit is set for the service group. ■ Type and dimension: integer-scalar ■ Default: Not applicable
OnlineRetryInterval (user-defined)	Indicates the interval, in seconds, during which a service group that has successfully restarted on the same system and faults again should be failed over, even if the attribute OnlineRetryLimit is non-zero. This prevents a group from continuously faulting and restarting on the same system. ■ Type and dimension: integer-scalar ■ Default: 0
OnlineRetryLimit (user-defined)	If non-zero, specifies the number of times the VCS engine tries to restart a faulted service group on the same system on which the group faulted, before it gives up and tries to fail over the group to another system. ■ Type and dimension: integer-scalar ■ Default: 0
OperatorGroups (user-defined)	List of operating system user groups that have Operator privileges on the service group. This attribute is valid only in secure clusters. ■ Type and dimension: string-keylist ■ Default: ""
Operators (user-defined)	List of VCS users with privileges to operate the group. A Group Operator can only perform online/offline, and temporary freeze/unfreeze operations pertaining to a specific group. See “About the VCS user privilege model” on page 63. ■ Type and dimension: string-keylist ■ Default: ""

Table C-3 Service group attributes

Service Group Attributes	Definition
Parallel (user-defined)	Indicates if service group is failover (0), parallel (1), or hybrid(2). ■ Type and dimension: integer-scalar ■ Default: 0
PathCount (system use only)	Number of resources in path not yet taken offline. When this number drops to zero, the engine may take the entire service group offline if critical fault has occurred. ■ Type and dimension: integer-scalar ■ Default: Not applicable
PreOnline (user-defined)	Indicates that the VCS engine should not online a service group in response to a manual group online, group autostart, or group failover. The engine should instead call a user-defined script that checks for external conditions before bringing the group online. ■ Type and dimension: boolean-scalar ■ Default: 0
PreOnlining (system use only)	Indicates that VCS engine invoked the preonline script; however, the script has not yet returned with group online. ■ Type and dimension: integer-scalar ■ Default: Not applicable
PreonlineTimeout (user-defined)	Defines the maximum amount of time in seconds the preonline script takes to run the command `hagrp -online -nopre` for the group. Note that HAD uses this timeout during evacuation only. For example, when a user runs the command `hastop -local -evacuate` and the Preonline trigger is invoked on the system on which the service groups are being evacuated. ■ Type and dimension: integer-scalar ■ Default: 300

Table C-3 Service group attributes

Service Group Attributes	Definition
Prerequisites (user-defined)	An unordered set of name=value pairs denoting specific resources required by a service group. If prerequisites are not met, the group cannot go online. The format for Prerequisites is: Prerequisites() = {Name=Value, name2=value2}. Names used in setting Prerequisites are arbitrary and not obtained from the system. Coordinate name=value pairs listed in Prerequisites with the same name=value pairs in Limits(). See "About system limits and service group prerequisites" on page 473. ■ Type and dimension: integer-association
PrintTree (user-defined)	Indicates whether or not the resource dependency tree is written to the configuration file. The value 1 indicates the tress is written. ■ Type and dimension: boolean-scalar ■ Default: 1
Priority (user-defined)	Enables users to designate and prioritize the service group. VCS does not interpret the value; rather, this attribute enables the user to configure the priority of a service group and the sequence of actions required in response to a particular event. ■ Type and dimension: integer-scalar ■ Default: 0
Probed (system use only)	Indicates whether all enabled resources in the group have been detected by their respective agents. ■ Type and dimension: boolean-scalar ■ Default: Not applicable
ProbesPending (system use only)	The number of resources that remain to be detected by the agent on each system. ■ Type and dimension: integer-scalar ■ Default: Not applicable
Responding (system use only)	Indicates VCS engine is responding to a failover event and is in the process of bringing the service group online or failing over the node. ■ Type and dimension: integer-scalar ■ Default: Not applicable

Table C-3 Service group attributes

Service Group Attributes	Definition
Restart (system use only)	For internal use only. ■ Type and dimension: integer-scalar ■ Default: Not applicable
SourceFile (system use only)	File from which the configuration was read. Always set to `./main.cf`. Make sure the path exists on all nodes before configuring this attribute. ■ Type and dimension: string-scalar ■ Default: ./main.cf
State (system use only)	Group state on each system: OFFLINE– All non-persistent resources are offline. ONLINE –All resources whose AutoStart attribute is equal to 1 are online. FAULTED–At least one critical resource in the group is faulted or is affected by a fault. PARTIAL–At least one, but not all, resources with Operations=OnOff is online, and not all AutoStart resources are online. STARTING–Group is attempting to go online. STOPPING– Group is attempting to go offline. A group state may be a combination of the multiple states described above. For example, OFFLINE | FAULTED, OFFLINE | STARTED, PARTIAL | FAULTED, PARTIAL | STARTING, PARTIAL | STOPPING, ONLINE | STOPPING ■ Type and dimension: integer-scalar ■ Default: Not applicable.
SystemList (user-defined)	List of systems on which the service group is configured to run and their priorities. Lower numbers indicate a preference for the system as a failover target. **Note:** You must define this attribute prior to setting the AutoStartList attribute. ■ Type and dimension: integer-association ■ Default: "" (none)

Table C-3 Service group attributes

Service Group Attributes	Definition
SystemZones (user-defined)	Indicates the virtual sublists within the SystemList attribute that grant priority in failing over. Values are string/integer pairs. The string key is the name of a system in the SystemList attribute, and the integer is the number of the zone. Systems with the same zone number are members of the same zone. If a service group faults on one system in a zone, it is granted priority to fail over to another system within the same zone, despite the policy granted by the FailOverPolicy attribute. ■ Type and dimension: integer-association ■ Default: "" (none)
Tag (user-defined)	Identifies special-purpose service groups created for specific VCS products. ■ Type and dimension: string-scalar ■ Default: Not applicable.
TargetCount (system use only)	Indicates the number of target systems on which the service group should be brought online. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
TFrozen (user-defined)	Indicates if service groups can be brought online on the system. Groups cannot be brought online if the attribute value is 1. ■ Type and dimension: boolean-scalar ■ Default: 0 (not frozen)
ToQ (system use only)	Indicates the node name to which the service is failing over. This attribute is specified when service group failover is a direct consequence of the group event, such as a resource fault within the group or a group switch. ■ Type and dimension: string-association ■ Default: Not applicable
TriggerEvent (system use only)	For internal use only. ■ Type and dimension: boolean-scalar ■ Default: Not applicable

Table C-3 Service group attributes

Service Group Attributes	Definition
TriggerResFault (user-defined)	Defines whether VCS invokes the resfault trigger when a resource faults. The value 0 indicates that VCS does not invoke the trigger. See "resfault event trigger" on page 535. ■ Type and dimension: integer-scalar ■ Default: 1
TriggerResStateChange (user-defined)	Determines whether or not to invoke the resstatechange trigger if resource state changes. See "resstatechange event trigger" on page 537. ■ Type and dimension: boolean-scalar ■ Default: 0 (disabled)
TypeDependencies (user-defined)	Creates a dependency (via an ordered list) between resource types specified in the service group list, and all instances of the respective resource type. ■ Type and dimension: string-keylist ■ Default: ""

Table C-3 Service group attributes

Service Group Attributes	Definition
UserIntGlobal (user-defined)	Use this attribute for any purpose. It is not used by VCS. ■ Type and dimension: integer-scalar ■ Default: 0
UserStrGlobal (user-defined)	VCS uses this attribute in the ClusterService group. Do not modify this attribute in the ClusterService group. Use the attribute for any purpose in other service groups. ■ Type and dimension: string-scalar ■ Default: 0
UserIntLocal (user-defined)	Use this attribute for any purpose. It is not used by VCS. ■ Type and dimension: integer-scalar ■ Default: 0
UserStrLocal (user-defined)	Use this attribute for any purpose. It is not used by VCS. ■ Type and dimension: string-scalar ■ Default: ""

System attributes

Table C-4 lists the system attributes.

Table C-4 System attributes

System Attributes	Definition
AgentsStopped (system use only)	This attribute is set to 1 on a system when all agents running on the system are stopped. ■ Type and dimension: integer-scalar ■ Default: Not applicable
AvailableCapacity (system use only)	Indicates system's available capacity when trigger is fired. If this value is negative, the argument contains the prefix % (percentage sign); for example, %-4. ■ Type and dimension: integer-scalar ■ Default: Not applicable
Capacity (user-defined)	Value expressing total system load capacity. This value is relative to other systems in the cluster and does not reflect any real value associated with a particular system. For example, the administrator may assign a value of 200 to a 16-processor machine and 100 to an 8-processor machine. ■ Type and dimension: integer-scalar ■ Default: 100
ConfigBlockCount (system use only)	Number of 512-byte blocks in configuration when the system joined the cluster. ■ Type and dimension: integer-scalar ■ Default: Not applicable
ConfigCheckSum (system use only)	Sixteen-bit checksum of configuration identifying when the system joined the cluster. ■ Type and dimension: integer-scalar ■ Default: Not applicable
ConfigDiskState (system use only)	State of configuration on the disk when the system joined the cluster. ■ Type and dimension: integer-scalar ■ Default: Not applicable
ConfigFile (user-defined)	Directory containing the configuration files. ■ Type and dimension: string-scalar ■ Default: ""

Table C-4 System attributes

System Attributes	Definition
ConfigInfoCnt (system use only)	The count of outstanding CONFIG_INFO messages the local node expects from a new membership message. This attribute is non-zero for the brief period during which new membership is processed. When the value returns to 0, the state of all nodes in the cluster is determined. ■ Type and dimension: integer-scalar ■ Default: Not applicable
ConfigModDate (system use only)	Last modification date of configuration when the system joined the cluster. ■ Type and dimension: integer-scalar ■ Default: Not applicable
CPUUsage (system use only)	Indicates the system's CPU usage by CPU percentage utilization. This attribute's value is valid if the Enabled value in the CPUUsageMonitoring attribute (below) equals 1. The value of this attribute is updated when there is a change of five percent since the last indicated value. ■ Type and dimension: integer-scalar ■ Default: Not applicable
CPUUsageMonitoring	Monitors the system's CPU usage using various factors. The values for ActionTimeLimit and NotifiyTimeLimt represent the time in seconds. The values for ActionThreshold and NotifyThreshold represent the threshold in terms of CPU percentage utilization. See "Monitoring CPU usage" on page 632. ■ Type and dimension: string-association ■ Default: Enabled = 0, NotifyThreshold = 0, NotifyTimeLimit = 0, ActionThreshold = 0, ActionTimeLimit = 0, Action = NONE.
CurrentLimits (system use only)	System-maintained calculation of current value of Limits. CurrentLimits = Limits - (additive value of all service group Prerequisites). ■ Type and dimension: integer-association ■ Default: Not applicable
DiskHbStatus (system use only)	Deprecated attribute. Indicates status of communication disks on any system. ■ Type and dimension: string-association ■ Default: Not applicable

Table C-4 System attributes

System Attributes	Definition
DynamicLoad (user-defined)	System-maintained value of current dynamic load. The value is set external to VCS with the `hasys -load` command. ■ Type and dimension: integer-scalar ■ Default: 0
EngineRestarted (system use only)	Indicates whether the VCS engine (HAD) was restarted by the hashadow process on a node in the cluster. The value 1 indicates that the engine was restarted; 0 indicates it was not restarted. ■ Type and dimension: boolean-scalar ■ Default: 0
Frozen (user-defined)	Indicates if service groups can be brought online on the system. Groups cannot be brought online if the attribute value is 1. ■ Type and dimension: boolean-scalar ■ Default: False
EngineVersion	Specifies the major, minor, maintenance-patch, and point-patch version of VCS. The value of EngineVersion attribute is in hexa-decimal format. To retrieve version information: `Major Version: EngineVersion >> 24 & 0xff Minor Version: EngineVersion >> 16 & 0xff Maint Patch: EngineVersion >> 8 & 0xff Point Patch : EngineVersion & 0xff` ■ Type and dimension: integer-scalar ■ Default: Not applicable
GUIIPAddr (user-defined)	Determines the local IP address that VCS uses to accept connections. Incoming connections over other IP addresses are dropped. If GUIIPAddr is not set, the default behavior is to accept external connections over all configured local IP addresses. See “User privileges for CLI commands” on page 66. ■ Type and dimension: string-scalar ■ Default: ““

Table C-4 System attributes

System Attributes	Definition
LicenseType (system use only)	Indicates the license type of the base VCS key used by the system. Possible values are: 0—DEMO 1—PERMANENT 2—PERMANENT_NODE_LOCK 3—DEMO_NODE_LOCK 4—NFR 5—DEMO_EXTENSION 6—NFR_NODE_LOCK 7—DEMO_EXTENSION_NODE_LOCK ■ Type and dimension: integer-scalar ■ Default: Not applicable
Limits (user-defined)	An unordered set of name=value pairs denoting specific resources available on a system. Names are arbitrary and are set by the administrator for any value. Names are not obtained from the system. The format for Limits is: Limits = { Name=Value, Name2=Value2}. ■ Type and dimension: integer-association ■ Default: ""
LinkHbStatus (system use only)	Indicates status of private network links on any system. ■ Type and dimension: string-association ■ Default: Not applicable
LLTNodeId (system use only)	Displays the node ID defined in the file /etc/llttab. ■ Type and dimension: integer-scalar ■ Default: Not applicable
LoadTimeCounter (system use only)	System-maintained internal counter of how many seconds the system load has been above LoadWarningLevel. This value resets to zero anytime system load drops below the value in LoadWarningLevel. ■ Type and dimension: integer-scalar ■ Default: Not applicable

Table C-4 System attributes

System Attributes	Definition
LoadTimeThreshold (user-defined)	How long the system load must remain at or above LoadWarningLevel before the LoadWarning trigger is fired. If set to 0 overload calculations are disabled. ■ Type and dimension: integer-scalar ■ Default: 600
LoadWarningLevel (user-defined)	A percentage of total capacity where load has reached a critical limit. If set to 0 overload calculations are disabled. For example, setting LoadWarningLevel = 80 sets the warning level to 80 percent. The value of this attribute can be set from 1 to 100. If set to 1, system load must equal 1 percent of system capacity to begin incrementing the LoadTimeCounter. If set to 100, system load must equal system capacity to increment the LoadTimeCounter. ■ Type and dimension: integer-scalar ■ Default: 80
NoAutoDisable (system use only)	When set to 0, this attribute autodisables service groups when the VCS engine is taken down. Groups remain autodisabled until the engine is brought up (regular membership). Setting this attribute to 1 bypasses the autodisable feature. ■ Type and dimension: boolean-scalar ■ Default: 0
NodeId (system use only)	System (node) identification specified in /etc/llttab. ■ Type and dimension: integer-scalar ■ Default: Not applicable
OnGrpCnt (system use only)	Number of groups that are online, or about to go online, on a system. ■ Type and dimension: integer-scalar ■ Default: Not applicable

Table C-4 System attributes

System Attributes	Definition
ShutdownTimeout (user-defined)	Determines whether to treat system reboot as a fault for service groups running on the system. On many systems, when a reboot occurs the processes are stopped first, then the system goes down. When the VCS engine is stopped, service groups that include the failed system in their SystemList attributes are autodisabled. However, if the system goes down within the number of seconds designated in ShutdownTimeout, service groups previously online on the failed system are treated as faulted and failed over. If you do not want to treat the system reboot as a fault, set the value for this attribute to 0. ■ Type and dimension: integer-scalar ■ Default: 120 seconds
SourceFile (user-defined)	File from which the configuration was read. Always set to ./main.cf. Make sure the path exists on all nodes before configuring this attribute. ■ Type and dimension: string-scalar ■ Default: ./main.cf
SysInfo (system use only)	Provides platform-specific information, including the name, version, and release of the operating system, the name of the system on which it is running, and the hardware type. ■ Type and dimension: string-scalar ■ Default: Not applicable
SysName (system use only)	Indicates the system name. ■ Type and dimension: string-scalar ■ Default: Not applicable
SysState (system use only)	Indicates system states, such as RUNNING, FAULTED, EXITED, etc. ■ Type and dimension: integer-scalar ■ Default: Not applicable
SystemLocation (user-defined)	Indicates the location of the system. ■ Type and dimension: string-scalar ■ Default: Not applicable

Table C-4 System attributes

System Attributes	Definition
SystemOwner (user-defined)	This attribute is used for VCS email notification and logging. VCS sends email notification to the person designated in this attribute when an event occurs related to the system. VCS also logs the owner name when an event occurs. ■ Type and dimension: string-scalar ■ Default: "unknown".
TFrozen (user-defined)	Indicates if a group can be brought online or taken offline. ■ Type and dimension: boolean-scalar ■ Default: 0
TRSE (system use only)	Indicates in seconds the time to Regular State Exit. Time is calculated as the duration between the events of VCS losing port h membership and of VCS losing port a membership of GAB. ■ Type and dimension: integer-scalar ■ Default: Not applicable
UpDownState (system use only)	This attribute has four values: 0 (down): System is powered off, or GAB and LLT are not running on the system. 1 (Up but not in cluster membership): GAB and LLT are running but the VCS engine is not. 2 (up and in jeopardy): The system is up and part of cluster membership, but only one network link (LLT) remains. 3 (up): The system is up and part of cluster membership, and has at least two links to the cluster. ■ Type and dimension: integer-scalar ■ Default: Not applicable
UserInt (user-defined)	Stores a system's integer value. ■ Type and dimension: integer-scalar ■ Default: 0
VCSFeatures (system use only)	Indicates which VCS features are enabled. Possible values are: 0—No features enabled (VCS Simulator) 1—L3+ is enabled 2—Global Cluster Option is enabled ■ Type and dimension: integer-scalar ■ Default: Not applicable

Cluster attributes

Table C-5 lists the cluster attributes.

Table C-5 Cluster attributes

Cluster Attributes	Definition
AdministratorGroups (user-defined)	List of operating system user account groups that have administrative privileges on the cluster. This attribute is valid only in secure clusters. ■ Type and dimension: string-keylist ■ Default: ""
Administrators (user-defined)	Contains list of users with Administrator privileges. ■ Type and dimension: string-keylist ■ Default: ""
AutoStartTimeout (user-defined)	If the local cluster cannot communicate with one or more remote clusters, this attribute specifies the number of seconds the VCS engine waits before initiating the AutoStart process for an AutoStart global service group. ■ Type and dimension: integer-scalar ■ Default: 150 seconds
BackupInterval (user-defined)	Time period in minutes after which VCS backs up the configuration files. The value 5 indicates VCS backs up configuration files every 5 minutes. You must set the configuration to read-write to enable backups. The value 0 indicates VCS does not back up configuration files. Set this attribute to at least 3. See “Scheduling automatic backups for VCS configuration files” on page 242. ■ Type and dimension: integer-scalar ■ Default: 0
ClusState (system use only)	Indicates the current state of the cluster. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
ClusterAddress (system use only)	Specifies the cluster’s virtual IP address (used by a remote cluster when connecting to the local cluster). ■ Type and dimension: string-scalar ■ Default: Not applicable.

Table C-5 Cluster attributes

Cluster Attributes	Definition
ClusterLocation (user-defined)	Specifies the location of the cluster. ■ Type and dimension: string-scalar ■ Default: ""
ClusterName (user-defined)	The name of cluster. ■ Type and dimension: string-scalar ■ Default: ""
ClusterOwner (user-defined)	This attribute is used for VCS notification; specifically, VCS sends notifications to persons designated in this attribute when an event occurs related to the cluster. See"About VCS event notification" on page 512. ■ Type and dimension: string-scalar ■ Default: "unknown"
ClusterTime (system use only)	The number of seconds since January 1, 1970. This is defined by the lowest node in running state. ■ Type and dimension: string-scalar ■ Default: Not applicable
ClusterUUID (system use only)	Unique ID assigned to the cluster by Availability Manager. ■ Type and dimension: string-scalar ■ Default: Not applicable
CompareRSM (user-defined)	Indicates if VCS engine is to verify that replicated state machine is consistent. This can be set by running the `hadebug` command. ■ Type and dimension: integer-scalar ■ Default: 0
ConnectorState (system use only)	Indicates the state of the wide-area connector (wac). If 0, wac is not running. If 1, wac is running and communicating with the VCS engine. ■ Type and dimension: integer-scalar ■ Default: Not applicable.

Table C-5 Cluster attributes

Cluster Attributes	Definition
CounterInterval (user-defined)	Intervals counted by the attribute GlobalCounter indicating approximately how often a broadcast occurs that will cause the GlobalCounter attribute to increase. The default value of the GlobalCounter increment can be modified by changing CounterInterval. If you increase this attribute to exceed five seconds, consider increasing the default value of the ShutdownTimeout attribute. ■ Type and dimension: integer-scalar ■ Default: 5
CredRenewFrequency	The number of days after which the VCS engine renews its credentials with the authentication broker. For example, the value 5 indicates that credentials are renewed every 5 days; the value 0 indicates that credentials are not renewed. ■ Type and dimension: integer-scalar ■ Default = 0
DumpingMembership (system use only)	Indicates that the engine is writing to disk. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
EnginePriority (user-defined)	The priority in which HAD runs. ■ Type and dimension: string-scalar ■ Default: Not applicable.
EngineShutdown (user-defined)	Defines the options for the hastop command. The attribute can assume the following values: Enable—Process all hastop commands. This is the default behavior. Disable—Reject all hastop commands. DisableClusStop—Do not process the hastop -all command; process all other hastop commands. PromptClusStop—Prompt for user confirmation before running the hastop -all command; process all other hastop commands. PromptLocal—Prompt for user confirmation before running the hastop -local command; reject all other hastop commands. PromptAlways—Prompt for user confirmation before running any hastop command. ■ Type and dimension: string-scalar ■ Default: Enable

Table C-5 Cluster attributes

Cluster Attributes	Definition
GlobalCounter (system use only)	This counter increases incrementally by one for each counter interval. It increases when the broadcast is received. VCS uses the GlobalCounter attribute to measure the time it takes to shut down a system. By default, the GlobalCounter attribute is updated every five seconds. This default value, combined with the 120-second default value of ShutdownTimeout, means if system goes down within twelve increments of GlobalCounter, it is treated as a fault. The default value of GlobalCounter increment can be modified by changing the CounterInterval attribute. ■ Type and dimension: integer-scalar ■ Default: Not applicable.
Guests (user-defined)	List of operating system user account groups that have Guest privileges on the cluster. This attribute is valid only in secure clusters. ■ Type and dimension: string-keylist ■ Default: ""
GroupLimit (user-defined)	Maximum number of service groups. ■ Type and dimension: integer-scalar ■ Default: 200
HacliUserLevel (user-defined)	This attribute has two, case-sensitive values: NONE–hacli is disabled for all users regardless of role. COMMANDROOT–hacli is enabled for root only. **Note:** The command `haclus -modify HacliUserLevel` can be executed by root only. ■ Type and dimension: string-scalar ■ Default: NONE
LockMemory (user-defined)	Controls the locking of VCS engine pages in memory. This attribute has the following values. Values are case-sensitive: ALL: Locks all current and future pages. CURRENT: Locks current pages. NONE: Does not lock any pages. ■ Type and dimension: string-scalar ■ Default: ALL

Table C-5 Cluster attributes

Cluster Attributes	Definition
LogSize (user-defined)	Size of engine log file. Minimum value = 64KB Maximum value = 128MB ■ Type and dimension: integer-scalar ■ Default: 33554432
Notifier (system use only)	Indicates the status of the notifier in the cluster; specifically: State—Current state of notifier, such as whether or not it is connected to VCS. Host—The host on which notifier is currently running or was last running. Default = None Severity—The severity level of messages queued by VCS for notifier. Values include Information, Warning, Error, and SevereError. Default = Warning Queue—The size of queue for messages queued by VCS for notifier. ■ Type and dimension: string-association ■ Default: Different values for each parameter.
OperatorGroups (user-defined)	List of operating system user groups that have Operator privileges on the cluster. This attribute is valid only in secure clusters. ■ Type and dimension: string-keylist ■ Default: ""
Operators (user-defined)	List of users with Cluster Operator privileges. ■ Type and dimension: string-keylist ■ Default: ""
PanicOnNoMem (system use only)	For internal use only. ■ Type and dimension: boolean-scalar ■ Default: Not applicable.
PrintMsg (user-defined)	Enables logging TagM messages in engine log if set to 1. ■ Type and dimension: boolean-scalar ■ Default: 0
ProcessPriority (user-defined)	The priority of HAD processes (for example, triggers). ■ Type and dimension: string-scalar ■ Default: "“

Table C-5 Cluster attributes

Cluster Attributes	Definition
ReadOnly (user-defined)	Indicates that cluster is in read-only mode. ■ Type and dimension: integer-scalar ■ Default: 1
ResourceLimit (user-defined)	Maximum number of resources. ■ Type and dimension: integer-scalar ■ Default: 5000
SecureClus	Indicates whether the cluster runs in secure mode. The value 1 indicated the cluster runs in secure mode. This attribute cannot be modified when VCS is running. ■ Type and dimension: boolean-scalar ■ Default: 0
SourceFile (user-defined)	File from which the configuration was read. Always set to ./main.cf. Make sure the path exists on all nodes before configuring this attribute. ■ Type and dimension: string-scalar ■ Default: Not applicable.
Stewards (user-defined)	The IP address and hostname of systems running the steward process. ■ Type and dimension: string-keylist ■ Default: Not applicable.
TypeLimit (user-defined)	Maximum number of resource types. ■ Type and dimension: integer-scalar ■ Default: 100
UserNames (user-defined)	List of VCS users. **Note:** Default user name is “admin”. ■ Type and dimension: string-association ■ Default: ""
VCSi3Info (system use only)	Enables VCS service groups to be mapped to Veritas i3 applications. This attribute is managed by the i3 product and should not be set or modified by the user. Contact your local Veritas sales representative for more information on the benefits of integrating VCS availability management with i3 ■ Type and dimension: string-association ■ Default: ""

Table C-5 Cluster attributes

Cluster Attributes	Definition
VCSFeatures (system use only)	Indicates which VCS features are enabled. Possible values are: 0—No features are enabled (VCS Simulator) 1—L3+ is enabled 2—Global Cluster Option is enabled ■ Type and dimension: integer-scalar ■ Default: Not applicable.
VCSMode (system use only)	Denotes the mode for which VCS is licensed. ■ Type and dimension: integer-scalar ■ Default:Not applicable
WACPort (user-defined)	The TCP port on which the wac (Wide-Area Connector) process on the local cluster listens for connection from remote clusters. Type and dimension: integer-scalar ■ Default: 14155

Heartbeat attributes

Table C-6 lists the heartbeat attributes. These attributes apply to global clusters.

Table C-6 Heartbeat attributes

Heartbeat Attributes	Definition
AgentState (system use only)	The state of the heartbeat agent. ■ Type and dimension: integer-scalar ■ Default: INIT
Arguments (user-defined)	List of arguments to be passed to the agent entry points. For the Icmp agent, this attribute can be the IP address of the remote cluster. ■ Type and dimension: string-vector ■ Default: ""
AYAInterval (user-defined)	The interval in seconds between two heartbeats. ■ Type and dimension: integer-scalar ■ Default: 60 seconds
AYARetryLimit (user-defined)	The maximum number of lost heartbeats before the agent reports that heartbeat to the cluster is down. ■ Type and dimension: integer-scalar ■ Default: 3
AYATimeout (user-defined)	The maximum time (in seconds) that the agent will wait for a heartbeat AYA entry point to return ALIVE or DOWN before being canceled. ■ Type and dimension: integer-scalar ■ Default: 300
CleanTimeOut (user-defined)	Number of seconds within which the Clean entry point must complete or be canceled. ■ Type and dimension: integer-scalar ■ Default: 300 seconds
ClusterList (user-defined)	List of remote clusters. ■ Type and dimension: string-keylist ■ Default: ""
InitTimeout (user-defined)	Number of seconds within which the Initialize entry point must complete or be canceled. ■ Type and dimension: integer-scalar ■ Default: 300 seconds

Table C-6 Heartbeat attributes

Heartbeat Attributes	Definition
LogDbg (user-defined)	The log level for the heartbeat. ■ Type and dimension: string-keylist ■ Default: ""
State	The state of the heartbeat. ■ Type and dimension: integer-scalar ■ Default: Not applicable
StartTimeout (user-defined)	Number of seconds within which the Start entry point must complete or be canceled. ■ Type and dimension: integer-scalar ■ Default: 300 seconds
StopTimeout (user-defined)	Number of seconds within which the Stop entry point must complete or be canceled without stopping the heartbeat. ■ Type and dimension: integer-scalar ■ Default: 300 seconds

Appendix D

Administering Symantec Web Server

- About Symantec Web Server
- Getting Started
- Configuring ports for VRTSweb
- Managing VRTSweb SSL certificates
- Configuring SMTP notification for VRTSweb
- Configuring logging for VRTSweb
- Modifying the maximum heap size for VRTSweb

About Symantec Web Server

Symantec Web Server (VRTSweb) is a Web Server component shared by various Symantec Web consoles, including Veritas Cluster Server and Veritas Volume Replicator

This document describes how to administer VRTSweb and provides instructions for common configuration tasks. Note that changes to the VRTSweb configuration apply to all Web consoles sharing the Web server.

The Web server is installed at the following path:

■ Windows: C:\Program Files\Veritas\VRTSweb.

To administer the VRTSweb from the command line, you must run commands from the following paths:

■ Windows: C:\Program Files\Veritas\VRTSweb\bin\

Getting Started

Connect to the Web server to start administering it.

To connect to Symantec Web Server

1 Access the Web server using the configured port number, for example, http://hostname:8181/.

2 Accept the certificate by Symantec.

■ To view and select the available Web consoles, click the **Home**.

■ To view and configure ports, SMTP recipients, SMTP servers, and logging, click **Configuration**.

3 Review the information in the **Configuration** tab:

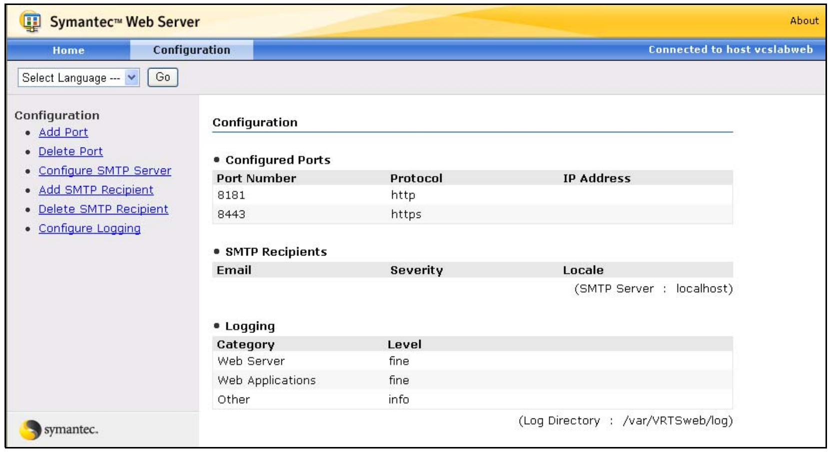

- Configured Ports—Lists information about the configured ports.
- SMTP Recipients—Displays information about configured SMTP recipients and the SMTP server.
- Logging—Lists the log levels for various Web server components.

Configuring ports for VRTSweb

You can view, add, edit, or remove ports for VRTSweb.

About VRTSweb ports

By default, VRTSweb serves HTML content on the following ports:

- 8181 (HTTP)—Non-secure port, used for backward compatibility. VRTSweb redirects users to the secure port
- 8443 (HTTPS)— Secure SSL port. VRTSweb presents a self-signed SSL certificate (issued by Symantec) to the browser. You must accept the certificate before accessing the secure Web consoles. The SSL protocol prevents malicious users from sniffing Web console data from the network.
- 14300—Administrative port.

If you use these ports for another application on the system, configure VRTSweb to use different ports.

Retrieving the list of VRTSweb ports

Retrieve the list of ports that are configured for VRTSweb.

To retrieve the list of ports from the command line

◆ On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
webgui listports
```

The output displays the list of configured ports and their protocols.

To retrieve the list of ports from the Web Console

1. Access the Web server using the configured port number, for example, http://hostname:8181/.
2. Click the **Configuration** tab.

   The **Configured Ports** table on the right side of the Configuration page lists the ports.

Adding ports for VRTSweb

Add ports for use by VRTSweb.

To add a port from the command line

◆ On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

webgui addport *portno protocol bind_ip_address*

- *portno*—The port number to be added.
- *protocol*—The protocol for the port. HTTP specifies a normal HTTP port, HTTPS specifies a secure SSL port.
  Web servers using the HTTP port can be accessed at http://*hostname:portno*/.
  Web servers using the HTTPS port can be accessed at https://*hostname:portno*/.
- *bind_ip_address*—Optional variable. Specifies that the new port be bound to a particular IP address instead of each IP address on the system. Use this option to restrict Web server access to specific administrative subnets.The IP address must be available on the system before the Web server starts.

For example:

addport 443 HTTPS .1.1.2

To add a port from the Web console

1 Access the Web server using the configured port number, for example, http://hostname:8181/.

2 Click the **Configuration** tab.

3 Click **Add Port** on the left side of the Configuration page.

4 In the Add Port pane:

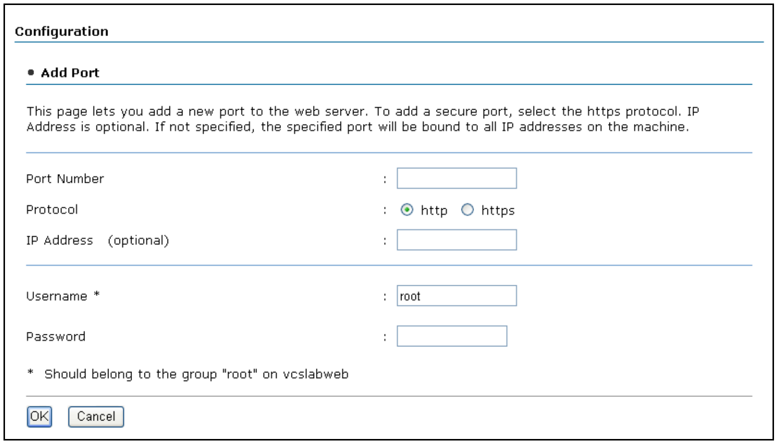

- Enter the port number to be added.
- Choose the HTTP option to add a normal port; choose the HTTPS option to add a secure SSL port.
  Web servers using the HTTP port can be accessed at http://*hostname:portno*/.
  Web servers using the HTTPS port can be accessed at https://*hostname:portno*/.
- Enter an IP address to bind the new port to a specific IP address instead of each IP address on the system. The IP address available on the system before starting the Web server. Use this attribute to restrict Web server access to specific administrative subnets.
- Enter the name and password for a user having superuser (administrative) privileges on the Web server system.

5 Click **OK**.

Deleting ports

Delete unused ports from VRTSweb.

To delete a port from the command line

- On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

  ```
  webgui delport <portno> [bind_ip_address]
  ```

  The variable *portno* represents the port number to be deleted. If the port was bound to a particular IP address, use the *bind_ip_address* option.

  You must ensure that at least one port remains configured for the Web server.

  For example:

  ```
  webgui delport 443 101.1.1.2
  webgui delport 80
  ```

To delete a port from the Web console

1. Access the Web server on a configured port. For example, http://hostname:8181/.
2. Click the **Configuration** tab.
3. Click **Delete Port** on the left side of the Configuration page.
4. In the Delete Port pane:
   - Enter the port number to be deleted. You cannot delete the port being used to access the Web page.
   - If the port was bound to a particular IP address, enter the IP address.
   - Enter the name and password for a user having superuser (administrative) privileges on Web server system.
5. Click **OK**.

Changing the administrative port

You can change the administrative port for VRTSweb only from the command line.

To change the administrative port

1. Stop the Web server:

   ```
   webgui stop force
   ```

2. Set the administrative port to a new value:

   ```
   webgui adminport new_port_no
   ```

3. Restart the Web server:

   ```
   webgui restart
   ```

Managing VRTSweb SSL certificates

VRTSweb presents a self-signed SSL certificate (issued by Symantec) when VRTSweb serves content over the secure port

Note: Certificate management commands are available only by the command line interface. Commands that modify the certificate require a server restart.

Viewing SSL certificate information

Display information about the configured SSL certificate.

To view information about the SSL certificate

◆ On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
webgui cert display
```

Creating a self-signed SSL certificate

Create a customized self-signed SSL certificate for VRTSweb.

To create a self-signed SSL certificate

1. Run the following interactive command on the system where VRTSweb is installed:
   ```
   webgui cert create
   ```
2. Follow the prompts to create a new certificate.
3. Restart the server for the new certificate to take effect.
   ```
   webgui restart
   ```

Exporting the SSL certificate to a file

You can export the public key that is associated with an SSL certificate to a file. You can then import the key to other applications to establish trust with the VRTSweb instance.

To export the SSL certificate to a file

- On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

  webgui cert export *cert_file* [rfc]

  If the VRTSweb SSL certificate does not exist, the command prompts you to create one. If you specify the RFC option, the key output is encoded in a printable format, which is defined by the Internet RFC 1421 standard.

  For example:

  ```
  %VRTSWEB_HOME%\bin> webgui cert export C:\myapp\vrtsweb.cer rfc
  ```

Configuring a CA-signed SSL certificate

By default, VRTSweb presents a self-signed SSL certificate every time you access VRTSweb over the SSL port. You can install a certificate signed by a Certificate Authority (CA) like Verisign.com or Thawte.com.

To configure a CA-signed SSL certificate

1. If you do not have a self-signed certificate with information that the CA can verify, create a certificate.

   webgui cert create

   See "Creating a self-signed SSL certificate" on page 752.

2. Generate a Certificate Signing Request (CSR) for the certificate. On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

   webgui cert certreq *certreq_file*

   The variable *certreq_file* specifies the file to which the CSR is written. The file is written using the Public-Key Cryptography Standard PKCS#10.

   For example:

   ```
   %VRTSWEB_HOME%\bin> webgui cert certreq c:\myapp\vrtsweb.cs
   ```

3. Submit the CSR to a certification authority, who issues a CA-signed certificate.

4. Import the CA-issued certificate to VRTSweb. On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

   webgui import *ca_cert_file*

The variable *cert_file* represents the certificate that is issued by the certification authority.

For example:

```
%VRTSWEB_HOME%\bin> webgui cert import c:\myapp\vrtsweb.cer
```

Note that the import command fails if the CA root certificate is not a part of the trust store that is associated with VRTSweb. If the command fails, add the CA root certificate to the VRTSweb trust store:

```
webgui cert trust ca_root_cert_file
```

For example:

```
%VRTSWEB_HOME%\bin> webgui cert trust c:\myapp\caroot.cer
```

Once the certificate used to sign the CSR is added to VRTSweb trust store, you can import the CA-assigned certificate into VRTSweb.

5 Restart VRTSweb:

```
webgui restart
```

Cloning the VRTSweb SSL certificate

You can clone the VRTSweb SSL keypair into a keystore and use the cloned VRTSweb certificate for another application or Web server. Visit http://java.sun.com for more information about keystores.

To clone the VRTSweb SSL certificate

- Run the following command:

```
webgui cert clone keystore storepass alias keypass
```

If a clone keystore exists, the command renames it to keystore.old. If the VRTSweb SSL certificate does not exist, the command prompts you to create one.

For example:

```
%VRTSWEB_HOME%\bin> webgui cert clone c:\myapp\myserv.keystore
        mystorepass myalias mykeypass
```

Configuring SMTP notification for VRTSweb

You can configure VRTSweb to send out email notifications about events that are associated with the Web server. For example:

- The Web server is starting/stopping [severity: INFORMATION]
- The Web console is starting/stopping [severity: INFORMATION]
- The Web server's allocated heap size very close to the maximum allowed [severity: SEVERE]

To send an email notification, VRTSweb needs to know the IP address or hostname of a configured SMTP server. The SMTP server address applies to all Web consoles running on the Web server. So you do not need to configure the SMTP server at multiple places.

Retrieving the name of the configured SMTP server

Retrieve the name of the SMTP server configured for VRTSweb.

To retrieve the name of the SMTP server from the command line

◆ On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
webgui smtp getserver
```

The command displays the SMTP server address or hostname, if it is configured.

To retrieve the name of the SMTP server from the Web console

1. Access the Web server on a configured port. For example, http://hostname:8181/
2. Click the **Configuration** tab.
3. The **SMTP Recipients** table on the right side of the page displays the configured SMTP server.

Setting the SMTP server

Configure an SMTP server for VRTSweb.

To set the SMTP server from the command line

◆ Run any of the following commands on the system where VRTSweb is installed:

```
webgui smtp setserver server_ip/hostname
webgui smtp delserver
```

The setserver command sets the SMTP server to the specified hostname/IP address. The delserver command deletes the SMTP server setting and disables SMTP notification.

For example:

```
%VRTSWEB_HOME%\bin> webgui smtp setserver
smtphost.company.com
%VRTSWEB_HOME%\bin> webgui smtp setserver 101.1.2.3
%VRTSWEB_HOME%\bin> webgui smtp delserver
```

To set the SMTP server from the Web Console

1. Access the Web server on a configured port. For example, http://hostname:8181/
2. Click the **Configuration** tab.
3. Click **Configure SMTP Server** on the left side of the Configuration page.
4. In the Configure SMTP Server dialog box:
   - Enter the IP address or hostname of the SMTP server to be used for notification. An empty string disables notification.
   - Enter the name and password for a user having superuser (administrative) privileges on the Web server system.
   - Click **OK**.

Retrieving SMTP settings

Retrieve configuration information for VRTSweb SMTP notification.

To retrieve SMTP recipients from the command line

◆ On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
webgui smtp listrcpt
```

This command retrieves the email addresses of the configured recipients, the notification severity level, and the notification locale.

To retrieve SMTP recipients from the Web console

1. Access the Web server on a configured port. For example, http://hostname:8181/
2. Click the **Configuration** tab.
3. The **SMTP Recipients** table on the right side of the Configuration page lists the configured SMTP recipients.

To retrieve the list of installed locales

◆ Run the following command:

```
webgui smtp listlocales
```

Adding an SMTP recipient

Add a user to receive SMTP notifications from VRTSweb.

To add an SMTP recipient from the command line

- On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
webgui smtp addrcpt email\
    [severity=<INFO|WARN|ERROR|SEVERE>] \
    [locale=<en|any_other_installed_locale>]
```

  - The variable *email* represents the email address of the new recipient.
  - The optional attribute *severity* represents the threshold for receiving Web server events. It can assume one of the following values: INFO|WARN|ERROR|SEVERE. If no value is specified for this attribute, it takes the default ERROR level.
  - The optional attribute *locale* specifies the locale in which the notification is to be sent. If no value is specified for this attribute, it takes the default locale of the system.

For example:

```
%VRTSWEB_HOME%\bin> webgui smtp addrcpt admin@company.com
severity=INFO locale=ja_JP
%VRTSWEB_HOME%\bin> webgui smtp addrcpt admin@company.com
severity=ERROR
%VRTSWEB_HOME%\bin> webgui smtp addrcpt admin@company.com
```

To add an SMTP recipient from the Web console

1. Access the Web server on a configured port. For example, http://hostname:8181/
2. Click the **Configuration** tab.
3. Click **Add SMTP Recipient** on the left side of the Configuration page.

4 In the Add SMTP Recipient dialog box:

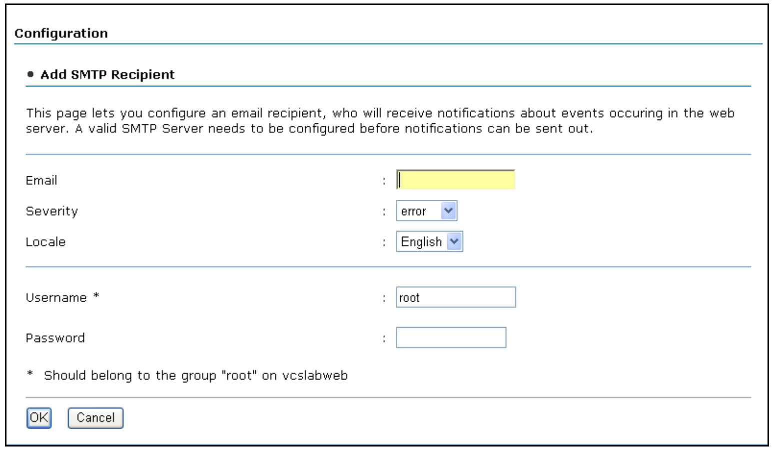

- **Email**—Email address of the new recipient.
- **Severity**—Threshold for receiving Web server events. Select one of the following values: INFO|WARN|ERROR|SEVERE.
- **Locale**—The locale in which notification is to be sent.
- **Username**—User having superuser (administrative) privileges on the Web server system
- **Password**—Password for the superuser.

5 Click **OK**.

Deleting an SMTP recipient

Delete an SMTP recipient to prevent VRTSweb from sending notifications to the recipient.

To delete an SMTP recipient from the command line

- On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

  ```
  webgui smtp delrcpt email
  ```

  The variable *email* represents the email address of the recipient to be deleted.

  For example:

  ```
  webgui smtp delrcpt admin@company.com
  ```

To delete an SMTP recipient from the Web console

1. Access the Web server on a configured port. For example, http://hostname:8181/
2. Click the **Configuration** tab.
3. Click **Delete SMTP Recipient** on the left side of the Configuration page.
4. In the Delete SMTP Recipient dialog box:
   - Enter the email address of the recipient to be deleted.
   - Enter the name and password for a user having superuser (administrative) privileges on the Web server system.
5. Click **OK.**

Configuring logging for VRTSweb

You can configure the amount of logs that individual VRTSWeb components generate. VRTSweb comprises the following components:

- Web server
- Web applications
- Other components

You can set the logging threshold for each component separately. The lower the threshold, the more are the logs generated. Symantec recommends setting log levels to lower values only for debugging.

Most of the logs are located at:

- %VRTSWEB_HOME%\log (for Windows),

Individual Symantec Web consoles choose their own locations for their logs. See the documentation for the specific Web console for more information.

Retrieving log levels

Display the current settings for VRTSweb logging.

To retrieve log levels from the command line

◆ On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
webgui log
```

This returns the logging thresholds for various components and the limit and rollover count of various log files for VRTSweb.

To retrieve log levels from the Web console

1. Access the Web server on a configured port. For example, http://hostname:8181/
2. Click the **Configuration** tab.
3. The **Logging** table on the right side of the Configuration page lists the log levels for various components of the Web server. The table does not display the limit and rollover count of various log files; you must use the command line to retrieve this information.

Modifying log levels for VRTSweb

Customize the log levels for VRTSweb.

To modify log levels from the command line

- On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

  webgui log [server=*level*] [webapps=*level*] [other=*level*]

  You can specify any of the following values for the variable *level* for each Web server component:
  FINE|FINER|FINEST|CONFIG|INFO|WARNING|SEVERE.

  Set the level to a lower value to generate more logs. FINEST is the lowest level while SEVERE is the highest level.

  For example:

  ```
  webgui log server=FINEST webapps=INFO Other=ERROR
  webgui log server=INFO
  ```

To modify log levels from rom the Web console

1 Access the Web server on a configured port. For example, http://hostname:8181/

2 Click the **Configuration** tab.

3 Click **Configure Logging** on the left side of the Configuration page.

4 In the Configure Logging dialog box:

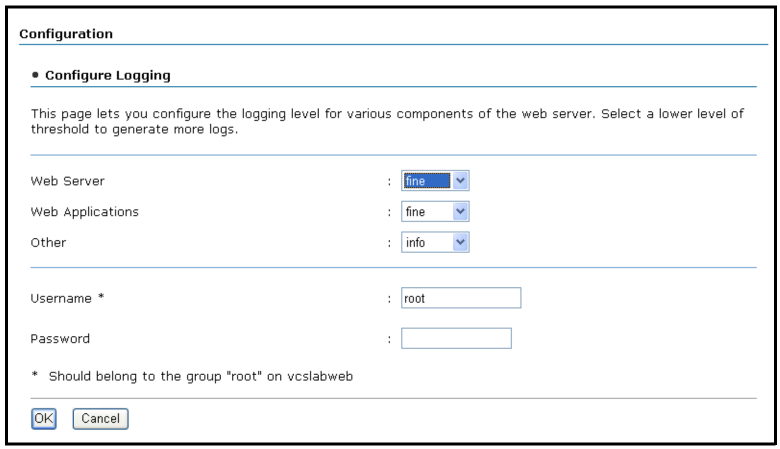

- Select the logging levels for the Web server, Web applications, and for other components.
- Enter the name and password for a user having superuser privileges on the Web server system.

5 Click **OK**.

Modifying size limit and rollover count for VRTSweb logs

You can modify the maximum size limit and rollover count for logs maintained by VRTSweb only from the command line.

To modify the size limit and rollover count for logs

- On the system where VRTSweb is installed, run the following command from the VRTSweb install directory:

```
$VRTSWEB_HOME/bin/webgui log
   [vrtsweb_size=size]          [vrtsweb_count=count]
   [command_size=size]          [command_count=count]
   [binary_size=size]           [binary_count=count]
   [jvm_size=size]              [jvm_count=count]
   [protocol_client_size=size]  [protocol_client_count=count]
   [protocol_server_size=size]  [protocol_server_count=count]
   [out_size=size]              [out_count=count]
   [err_size=size]              [err_count=count]
   [webapps_size=size]          [webapps_count=count]
```

  For example:

```
webgui log vrtsweb_size=100000 vrtsweb_count=4
webgui log err_size=200000
webgui log webapps_count=4
```

The following table describes the command parameters:

Parameter	Description
vrtsweb_size	The size of the file _vrtsweb.log, which contains the Web server logs and the tomcat container related logs.
vrtsweb_count	The count for the file _vrtsweb.log.
command_size	The size of the file _command.log, which contains the logs related to administrative commands.
command_count	The count for the file _command.log.
binary_size	The size of the file _binary.log, which contains the binary representation of other log files.
binary_count	The count for the file _binary.log.
jvm_size	The size of the file _jvm.log, which contains JVM-related measurements. The file records the memory that is consumed by the JVM at various times.
jvm_count	The count for the file _jvm.log.

Parameter	Description
protocol_client_size	The size of the file _protocol_client.log, which contains the communication sent (and received) between various utilities and the server.
protocol_client_count	The count for the file _protocol_client.log.
protocol_server_size	The size of the file _protocol_server.log, which contains the communication sent (and received) by the server to various utilities.
protocol_server_count	The count for the file _protocol_server.log.
out_size	The size of the file _out.log, which contains messages that are logged to the standard output stream of the JVM.
out_count	The count for the file _out.log.
err_size	The size of the file _err.log, which contains messages that are logged to the standard error stream of the JVM, including any stack traces.
err_count	The count for the file _err.log.
webapps_size	The default size for log files of all Web applications running VRTSweb. Individual Web applications can override this default value.
webapps_count	The count for log files of all Web applications running VRTSweb. Individual Web applications can override this default value.

Modifying the maximum heap size for VRTSweb

The default maximum allowed heap size for the VRTSWeb Java Virtual Machine (JVM) is 256MB. You may need to modify this limit for large configurations or for when a large number of consoles share the same VRTSweb instance.

You can modify the maximum heap size only from the command line.

To modify the maximum heap size

1. Type the following command:

   `webgui maxheap` ***`new_size_in_MB`***

   For example:

   ```
   webgui maxheap 512
   ```

2. Restart the Web server after specifying a new limit.

   ```
   webgui restart
   ```

3. View the current limit.

   ```
   webgui maxheap
   ```

Accessibility and VCS

- About accessibility in VCS
- Navigation and keyboard shortcuts
- Support for accessibility settings
- Support for assistive technologies

About accessibility in VCS

Veritas Cluster Server provides shortcuts for major graphical user interface (GUI) operations and menu items. Veritas Cluster Server is compatible with operating system accessibility settings as well as a variety of assistive technologies. All manuals also are provided as accessible PDF files, and the online help is provided as HTML displayed in a compliant viewer.

Navigation and keyboard shortcuts

VCS uses standard operating system navigation keys and keyboard shortcuts. For its unique functions, VCS uses its own navigation keys and keyboard shortcuts which are documented below.

Navigation in the Java Console

The following table lists keyboard navigation rules and shortcuts used in Cluster Manager (Java Console), in addition to those provided by the operating system:

VCS Keyboard Input	Result
[Shift F10]	Opens a context-sensitive pop-up menu
[Spacebar]	Selects an item
[Ctrl Tab]	Navigates outside a table
[F2]	Enables editing a cell

Navigation in the Web Console

Cluster Management Console supports standard browser-based navigation and shortcut keys for supported browsers.

All Symantec GUIs use the following keyboard navigation standards:

- Tab moves the focus to the next active area, field, or control, following a preset sequence. Shift+Tab moves the focus in the reverse direction through the sequence.
- Ctrl+Tab exits any Console area that you internally navigate with Tab.
- Up and Down arrow keys move focus up and down the items of a list.
- Alt in combination with the underlined mnemonic letter for a field or command button shifts the focus to that field or button.

- Either Enter or the Spacebar activates your selection. For example, after pressing Tab to select Next in a wizard panel, press the Spacebar to display the next screen.

Support for accessibility settings

Symantec software responds to operating system accessibility settings. Symantec products are compatible with accessibility utilities provided by operating systems.

On UNIX systems, you can change the accessibility settings using desktop preferences or desktop controls. On Windows 2000 systems, you can set accessibility options involving keyboard responsiveness, display contrast, alert sounds, and mouse operation through the Control Panel (**Start > Settings > Control Panel > Accessibility Options**) and through the Accessibility Wizard (**Start > Programs > Accessories > Accessibility > Accessibility Wizard**).

Support for assistive technologies

- Cluster Manager (Java Console) is compatible with JAWS 4.5.
- Though graphics in the documentation can be read by screen readers, setting your screen reader to ignore graphics may improve performance.
- Symantec has not tested screen readers for languages other than English.

Configuring LLT over UDP

VCS provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

Note: LLT over UDP is not supported on IPV6.

When to use LLT over UDP

Use LLT over UDP when:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Because LLT over UDP is slower than LLT over Ethernet, LLT over UDP should only be used when the hardware configuration makes it necessary.

Configuring LLT over UDP

The following is a checklist for configuring LLT over UDP. Examples are provided in the sections that follow.

- Make sure that each NIC has an IP address configured before configuring LLT. Each link must be in a different subnet. See the examples in the following sections.
- Make sure that each link has a unique UDP port; do not assign well-known ports. See "Selecting UDP ports" on page 774.
- Set the broadcast address correctly for direct-attached (non-routed) links.
- For links that cross an IP router, disable broadcast features and specify the IP address of each link manually in the `llttab` file. The default path for the

file is `%VCS_HOME%\comms\llt\llttab.txt` See “Sample configuration: Links crossing IP routers” on page 776.

The link command in the llttab file

The following table describes the fields of the `link` command shown in the `llttab` file examples that follow; see “Sample configuration: Direct-attached links” on page 775, and “Sample configuration: Links crossing IP routers” on page 776. Note that some of these fields differ from the command for standard LLT links.

`<tag-name>`	A unique string that is used as a tag by LLT; for example `link1`, `link2`,
`<device>`	The device path of the UDP protocol; for example `udp`
`<node-range>`	Nodes using the link. “`-`” indicates *all* cluster nodes are to be configured for this link.
`<link-type>`	Type of link; must be “`udp`” for LLT over UDP
`<udp-port>`	Unique UDP port in range of 49152-65535 for the link; see “Selecting UDP ports” on page 774.
`<MTU>`	“`-`” is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. The `lltstat -l` command can display the current value.
`<IP address>`	IP address of the link on the local node.
`<bcast-address>`	■ for clusters having broadcasts enabled, specify the value of the subnet broadcast address ■ “-” is the default for clusters spanning routers

The set-addr command in the llttab file

The `set-addr` command in the `llttab` file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers. The following table describes the fields of the `set-addr` command; see "Sample configuration: Links crossing IP routers" on page 776.

`<node-id>`	The ID of the cluster node; for example, 0.
`<link tag-name>`	The string used by LLT to identify the link; for example `link1, link2,..`
`<address>`	IP address assigned to the link on the peer node.

Selecting UDP ports

When selecting a UDP port, select an available 16-bit integer from the range described below.

- Use available ports (that is, ports that are not in use)] in the private range 49152 to 65535
- Do not use:
  - Ports from the range of well-known ports, 0 to 1023
  - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `C:\WINNT\system32\drivers\etc>services`. You should also use the `netstat` command to list the ports currently in use. For example:

```
# netstat -a -p UDP

Proto   Local Address          Foreign Address        State
  UDP    THORPC111:snmp          *:*
  UDP    THORPC111:snmptrap      *:*
  UDP    THORPC111:microsoft-ds   *:*
  UDP    THORPC111:isakmp        *:*
  UDP    THORPC111:1027          *:*
  UDP    THORPC111:1028          *:*
  UDP    THORPC111:1029          *:*
  UDP    THORPC111:1030          *:*
  UDP    THORPC111:1059          *:*
  UDP    THORPC111:1063          *:*
 UDP    THORPC111:4219          *:*
  UDP    THORPC111:4500          *:*
  UDP    THORPC111:ntp           *:*
  UDP    THORPC111:netbios-ns    *:*
  UDP    THORPC111:netbios-dgm   *:*
  UDP    THORPC111:ntp           *:*
  UDP    THORPC111:1646          *:*
  UDP    THORPC111:3217          *:*
  UDP    THORPC111:3219          *:*
  UDP    THORPC111:3456          *:*
```

Look in the UDP section of the output; UDP ports listed under `Local Address` are already in use. If a port is listed in the `services` file, its associated name is displayed rather than the port number in the output of the `netstat` command.

Sample configuration: Direct-attached links

The following illustration depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure F-6 Direct-attached links employing LLT over UDP

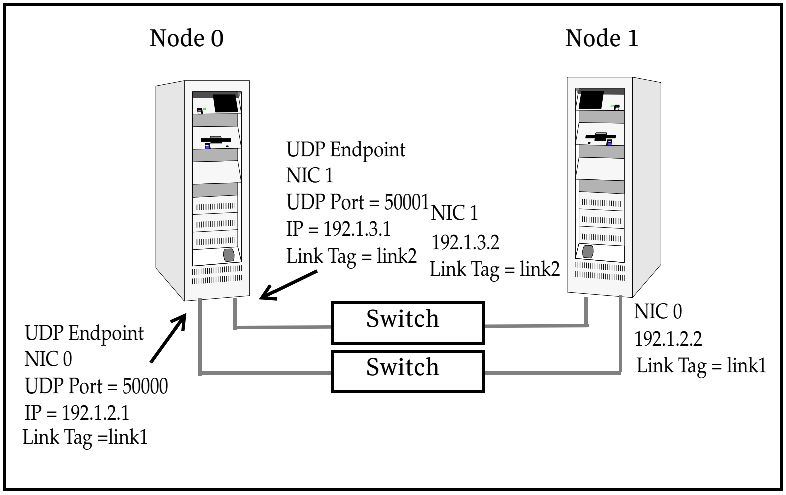

The configuration represented by the following `llttab` file for Node 0 has directly attached crossover links or links connected through a hub or switch. These links do not cross routers.

Because LLT broadcasts requests to peer nodes to discover their addresses, the addresses of peer nodes do not need to be specified in the `llttab` file using the `set-addr` command. For direct attached links, you need to set the broadcast address of the links in the `llttab` file. Verify that the IP addresses and broadcast addresses are set correctly.

```
set-node Node0
set-cluster 1
#configure Links
#link <tag-name> <device> <node-range> <link-type> <udp port>
<MTU> <IP-address> <bcast-address>
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 would resemble:

```
set-node Node1
set-cluster 1
#configure Links
#link <tag-name> <device> <node-range> <link-type> <udp port>
<MTU> <IP-address> <bcast-address>
link link1 udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: Links crossing IP routers

The following illustration depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows just two nodes of a four-node cluster.

Figure F-7 Links crossing an IP router employing LLT over UDP

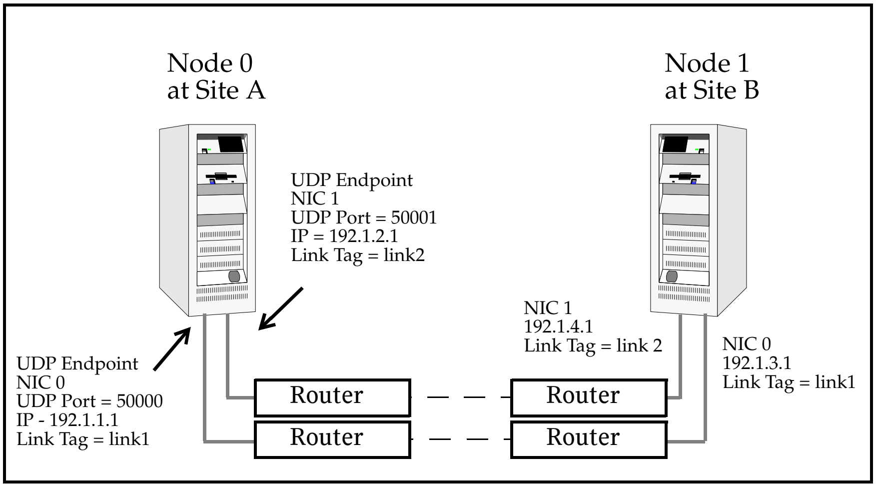

The configuration represented by the following `llttab` file for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. The broadcast features are disabled because LLT is unable to broadcast requests for addresses across routers, so the broadcast address does not need to be set in the in the `link` command of the `llttab` file.

```
set-node Node1
set-cluster 1

link link1 udp - udp 50000 - 192.1.3.1 -
link link2udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr <node-id> <link tag-name> <address>
set-addr        0 link1 192.1.1.1
set-addr        0 link2 192.1.2.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

The `lltab` file on Node 0 would resemble:

```
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 192.1.1.1 -
link link2 udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr <node-id> <link tag-name> <address>
set-addr        1 link1 192.1.3.1
set-addr        1 link2 192.1.4.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

Handling concurrency violation in Any-to-Any configurations

This appendix describes how you can use the Process agent, along with a sample script, to handle potential concurrency violation in an Any-to Any configuration that uses the GenericService agent.

Concurrency violation scenario

Consider the following excerpt from a configuration in which an antivirus software is configured as a GenericService resource in an Exchange 2000 service group.

Exchange Service Group 1

```
System List = {S1, S2}
GenericService G1 controlling "NortonAntivirus"
E2KServices   E1 managing Exchange Information Store, Message
                                                Transfer Agent,
and System Attendant services
Lanman E2KVName1 controlling the E2K Virtual server
...

G1 requires E1
E1 requires E2KVName1
...
```

Exchange Service Group 2:

```
System List = {S1, S3}
GenericService G2 controlling "NortonAntivirus"
E2KServices   E2 managing Exchange Information Store, Message
                                                Transfer Agent,
and System Attendant services
Lanman E2KVName2 controlling the E2K Virtual server
```

```
...

G1 requires E1
E1 requires E2KVName2
...
```

In this example, S3 is the standby system for both S1 and S2. Service Group 1 is online on S1 and Service Group 2 is online on S2. Such a configuration is desirable for an application like Exchange 2000 that requires an antivirus service instance attached to it.

Consider a scenario where Service Group 1 fails over from S1 to S3. When G1 comes online on S3, G2 also reports online because both G1 and G2 monitor the same service. As a result, VCS reports a concurrency violation for Service Group 2 from S3 and tries to take G2 offline on S3. As soon as G2 is taken offline on S3, G1 reports a fault, and Service Group 1 faults on S3.

This situation can be addressed by using a custom script along with the Process agent, in place of the GenericService agent.

About the vcsgensvc.vbs script

The script `vcsgensvc.vbs` resides at the path `%VCS_HOME%\Samples\Process`. The script works with the Process agent to bring services online, monitor them, and take them offline.

The script takes the following parameters:

Table G-7 vcsgensvc.vbs script parameters

Parameter	Accepted Values
Operation	online offline monitor
Service Name	Display or key name of the service
Computer Name	The name of the computer (virtual computer being monitored This parameter applies only to the MonitorProgram attribute of the Process resource.

Sample configuration to handle concurrency violation

In the following sample configuration, the Process agent monitors the Norton Antivirus service. The script is installed in the following directory:

```
D:\Program Files\Veritas\Cluster Server\Samples\Process
```

The script takes the Exchange virtual server name as an input parameter and monitors the service using this virtual name. If the Exchange Virtual server is online and the antivirus service is running, the script returns ONLINE, instead of returning ONLINE based on the status of the service alone.

Note that this recommendation is for this specific scenario only. The Process agent is not an alternative to the GenericService agent, which offers added functionality for generic services.

For Exchange Service Group 1:

```
System List = {S1, S2}
Process                    AVService1 controlling
"NortonAntivirus"
    E2KServices         E1
Lanman         E2KVName1 controlling the E2K Virtual server
...

G1 requires E1
E1 requires E2KVName1
...

Lanman L1 (
        VirtualName = E2KVName1
)

Process AVService1 (
        StartProgram = "CScript.exe \"d:\\program files\\
                       veritas\\cluster
server\\samples\\process\\
                       vcsgensvc.vbs\" online NortonAntivirus"
        StopProgram = "CScript.exe \"d:\\program files\\
                       veritas\\cluster
server\\samples\\process\\
                       vcsgensvc.vbs\" offline NortonAntivirus"
        MonitorProgram = "CScript.exe \"d:\\program files\\
                       veritas\\cluster
server\\samples\\process\\
                       vcsgensvc.vbs\" monitor NortonAntivirus
E2KVName1"
)
```

For Exchange Service Group 2:

```
System List = {S1, S3}
Process             AVService2 controlling "NortonAntivirus"
```

```
    E2KServices                        E2
Lanman         E2KVName2 controlling the E2K Virtual server
...

G1 requires E1
E1 requires E2KVName2
...

Lanman L2 (
        VirtualName = E2KVName2
)

Process AVService2 (
        StartProgram = "CScript.exe \"d:\\program
files\\veritas\\
                      cluster
server\\samples\\process\\vcsgensvc.vbs\"
                      online NortonAntivirus"
        StopProgram = "CScript.exe \"d:\\program
files\\veritas\\
                      cluster
server\\samples\\process\\vcsgensvc.vbs\"
                      offline NortonAntivirus"
        MonitorProgram = "CScript.exe \"d:\\program files\\
                      veritas\\cluster
server\\samples\\process\\
                      vcsgensvc.vbs\" monitor NortonAntivirus
E2KVName2"
)
```

Notes for using scripts with the Process agent

- In the above example, the supplied script assumes that Service Group 1 and Service Group 2 will never come online on one system. Service Group Workload Management or triggers must be configured to meet this requirement.
- When using this configuration, we recommend setting the user context of the Process agent to LocalSystem.

Glossary

Agent

A process that starts, stops, and monitors all configured resources of a type, and reports their status to VCS.

Active/Active Configuration

A failover configuration where each systems runs a service group. If either fails, the other one takes over and runs both service groups. Also known as a symmetric configuration.

Active/Passive Configuration

A failover configuration consisting of one service group on a primary system, and one dedicated backup system. Also known as an asymmetric configuration.

Authentication Broker

The Veritas Security Services component that serves, one level beneath the root broker, as an intermediate registration authority and a certification authority. The authentication broker can authenticate clients, such as users or services, and grant them a certificate that will become part of the Veritas credential. An authentication broker cannot, however, authenticate other brokers. That task must be performed by the root broker.

See "Root Broker."

Cluster

One or more computers linked together for the purpose of multiprocessing and high availability. The term is used synonymously with VCS cluster, meaning one or more computers that are part of the same GAB membership.

Daemon Down Node Alive (DDNA)

A situation where the VCS high availability daemon has failed on a system and has not been restarted by the hashadow process.

Disaster Recovery

A solution that supports fail over to a cluster in a remote location in the event that the local cluster becomes unavailable. Disaster recovery global clustering, heartbeating, and replication.

Failover

A failover occurs when a service group faults and is migrated to another system.

GAB

Group Atomic Broadcast (GAB) is a communication mechanism of the VCS engine that manages cluster membership, monitors heartbeat communication, and distributes information throughout the cluster.

Global Service Group

A VCS service group that spans across two or more clusters. The ClusterList attribute for the group contains the list of clusters over which the group spans.

hashadow Process

A process that monitors and, when required, restarts HAD.

High Availability Daemon (HAD)

The core VCS process that runs on each system. The HAD process maintains and communicates information about the resources running on the local system and receives information about resources running on other systems in the cluster.

Jeopardy

A node is in *jeopardy* when it is missing one of the two required heartbeat connections. When a node is running with one heartbeat only (in jeopardy), VCS does *not* restart the applications on a new node. This action of disabling failover is a safety mechanism that prevents data corruption.

LLT

Low Latency Transport (LLT) is a communication mechanism of the VCS engine that provides kernel-to-kernel communications and monitors network communications.

main.cf

The file in which the cluster configuration is stored.

Network Partition

If all network connections between any two groups of systems fail simultaneously, a *network partition* occurs. When this happens, systems on both sides of the partition can restart applications from the other side resulting in duplicate services, or “split-brain.” A split-brain occurs when two independent systems configured in a cluster assume they have exclusive access to a given resource (usually a file system or volume). The most serious problem caused by a network partition is that it affects the data on shared disks.

See “Jeopardy.”

See “Seeding.”

Node

The physical host or system on which applications and service groups reside. When systems are linked by VCS, they becomes nodes in a cluster.

N-to-1

An N-to-1 configuration is based on the concept that multiple, simultaneous server failures are unlikely; therefore, a single backup server can protect multiple active servers. When a server fails, its applications move to the backup server. For example, in a 4-to-1 configuration, one server can protect four servers, which reduces redundancy cost at the server level from 100 percent to 25 percent.

N-to-N

N-to-N refers to multiple service groups running on multiple servers, with each service group capable of being failed over to different servers in the cluster. For example, consider a four-node cluster with each node supporting three critical database instances. If any node fails, each instance is started on a different node, ensuring no single node becomes overloaded.

N-to-M

N-to-M (or Any-to-Any) refers to multiple service groups running on multiple servers, with each service group capable of being failed over to different servers in the same cluster, and also to different servers in a linked cluster. For example, consider a four-node cluster with each node supporting three critical database instances and a linked two-node back-up cluster. If all nodes in the four-node cluster fail, each instance is started on a node in the linked back-up cluster.

Replication

Replication is the synchronization of data between systems where shared storage is not feasible. The systems that are copied may be in local backup clusters or remote failover sites. The major advantage of replication, when compared to traditional backup methods, is that current data is continuously available.

Resources

Individual components that work together to provide application services to the public network. A resource may be a physical component such as a disk or network interface card, a software component such as Oracle8*i* or a Web server, or a configuration component such as an IP address or mounted file system.

Resource Dependency

A dependency between resources is indicated by the keyword "requires" between two resource names. This indicates the second resource (the child) must be online before the first resource (the parent) can be brought online. Conversely, the parent must be offline before the child can be taken offline. Also, faults of the children are propagated to the parent.

Resource Types

Each resource in a cluster is identified by a unique name and classified according to its type. VCS includes a set of predefined resource types for storage, networking, and application services.

Root Broker

The first authentication broker, which has a self-signed certificate. The root broker has a single private domain that holds only the names of brokers that shall be considered valid. See"Authentication Broker."

Seeding

Seeding is used to protect a cluster from a pre-existing network partition. By default, when a system comes up, it is not seeded. Systems can be seeded automatically or manually. Only systems that have been seeded can run VCS. Systems are seeded automatically only when: an unseeded system communicates with a seeded system or all systems in the cluster are unseeded and able to communicate with each other.

See "Network Partition."

Service Group

A service group is a collection of resources working together to provide application services to clients. It typically includes multiple resources, hardware- and software-based, working together to provide a single service.

Service Group Dependency

A mechanism by which two service groups can be linked by a dependency rule.

Shared Storage

Storage devices that are connected to and used by two or more systems.

SNMP Notification

Simple Network Management Protocol (SNMP) developed to manage nodes on an IP network.

State

The current activity status of a resource, group or system.

types.cf

The types.cf file describes standard resource types to the VCS engine; specifically, the data required to control a specific resource.

Virtual IP Address

A unique IP address that associated with the cluster. It may be brought up on any system in the cluster, along with the other resources of the service group. This address, also known as the IP alias should not be confused with the base IP address, which is the IP address that corresponds to the host name of a system.

Index

A

About 567
accessibility
    assistive technology support 769
    overview 767
ActionTimeout attribute 708
ActiveCount attribute 715
AdministratorGroups attribute
    for clusters 735
    for service groups 715
Administrators attribute
    for clusters 735
    for service groups 715
agent log
    format 637
    location 637
AgentClass attribute 708
AgentDirectory attribute 708
AgentFailedOn attribute 708
AgentFile attribute 708
AgentPriority attribute 708
AgentReplyTimeout attribute 709
agents
    classifications of 26
    entry points 25
    framework 26
    Heartbeat 547
    impact on performance 623
    starting from command line 264
    stopping from command line 264
    Wide-Area Heartbeat 547
AgentStartTimeout attribute 709
AgentState attribute 742
AgentStopped attribute 728
aggregate notifications, monitoring 525
alerts
    deleting from Java Console 233
    monitoring from Java Console 232
    types of 669
Application Configuration wizard 343
applications
    configuring in VCS 277
    configuring using wizard 343
ArgList attribute 709
ArgListValues attribute 702
assistive technology support 769
association attribute dimension 54
asymmetric configuration 34
AttrChangedTimeout attribute 709
attribute dimensions
    association 54
    keylist 54
    scalar 54
    vector 54
attribute types
    boolean 53
    integer 53
    string 53
attributes
    about 52
    editing from Cluster Management Console 131
    editing from Java Console 223
    for clusters 735
    for heartbeats 742
    for resource types 708
    for resources 702
    for service groups 715
    for systems 728
    local and global 56
    overriding from command line 274
    overriding from Java Console 209
authentication broker 29
Authority attribute
    about 547
    definition 715
AutoDisabled attribute 715
AutoFailover attribute
    about 451
    definition 716
AutoRestart attribute 716
AutoStart attribute
    for resources 702
    for service groups 716

AutoStartIfPartial attribute 716
AutoStartList attribute 716
AutoStartPolicy attribute 717
AutoStartTimeout attribute 735
AvailableCapacity attribute 728

B

BackupInterval attribute 735
binary message catalogs
  about 639
  location of 639
boolean attribute type 53
bundled agents 26
  configuring 277

C

Capacity attribute 728
Changing 447
Chapter 13, 61, 417, 509
CleanTimeout attribute 709
client process, detecting failure 630
CloseTimeout attribute 709
ClusState attribute 735
Cluster Administrator
  about 64
  adding user as 249
cluster attributes 735
Cluster Explorer
  about 146
  accessing 146
  adding resources 200
  adding service groups 178
  adding systems 219
  adding users 173
  autoenabling service groups 191
  bringing resources online 205
  bringing service groups online 183
  changing user passwords 175
  changing user privileges 176
  clearing resource faults 212
  clearing ResourceInfo attribute 218
  closing configuration files 223
  Cluster Query 164
  Command Center 162
  configuration tree 149
  deleting resources 205
  deleting service groups 182
  deleting users 174
  disabling resources 211
  disabling service groups 190
  editing attributes 223
  enabling resources 210
  enabling service groups 189
  flushing service groups 192
  freezing service groups 187
  freezing systems 221
  importing resource types 218
  linking resources 213
  linking service groups 193
  logs 230
  modifying system lists for service groups 160
  monitoring group dependencies 152
  monitoring resource dependencies 153
  Notifier Wizard 163
  opening configuration files 222
  probing resources 208
  Properties view 151
  refreshing ResourceInfo attribute 218
  Remote Cluster Status View 158
  Resource View 153
  saving configuration files 222
  service group configuration wizard 197
  Service Group View 152
  Status View 150
  switching service groups 186
  System Connectivity View 157
  System Manager 160
  taking resources offline 206
  taking resources offline and propagating 207
  taking service groups offline 185
  tear-off view 149
  Template View 159
  toolbar 147
  unfreezing service groups 188
  unfreezing systems 221
  unlinking resources 215
  unlinking service groups 195
  User Manager 161
  view panel 149
Cluster Guest
  about 65
  adding user as 249
Cluster Management Console
  administering resources 118, 129
  administering service groups 102
  managing cluster configuration 100
Cluster Manager (Java Console). See Java Console

Cluster Monitor
  about 140
  adding clusters 168
  administering 168
  behavior during failover 142
  collapsing displays 143
  configuring existing panels 169
  configuring new panels 168
  expanding displays 143
  icon colors 142
  logging off of a cluster 172
  logging on to a cluster 170
  menus 141
  monitoring cluster connection 142
  monitoring cluster objects 142
  panels 142
  pausing scrolling panels 143
  toolbar 141
cluster name, changing in global configuration 613
Cluster Operator
  about 65
  adding user as 249
Cluster Query
  in Java Console 164
ClusterAddress attribute 735
ClusterFailOverPolicy attribute 717
clustering
  criteria for data storage 19
  criteria for monitor procedure 18
  criteria for start procedure 18
  criteria for stop procedure 18
  license and host name issues 19
ClusterList attribute 717
ClusterLocation attribute 736
ClusterName attribute 736
ClusterOwner attribute 736
clusters
  adding nodes 372
  administering from Java Console 222
  connecting to Cluster Monitor 168
  deleting 401
  reconfiguring 384
  removing nodes 380
  setting up 73
ClusterService group
  configuring using the wizard 392
ClusterTime attribute 736
ClusterUUID attribute 736
Command Center
  accessing 162
  adding resources 201
  adding service groups 179
  adding systems 219
  autoenabling service groups 191
  bringing resources online 205
  bringing service groups online 184
  clearing resource faults 212
  closing configuration files 223
  deleting resources 205
  deleting service groups 182
  deleting systems 220
  disabling resources 211
  disabling service groups 190
  editing attributes 224
  enabling resources 210
  enabling service groups 189
  executing commands 223
  flushing service groups 192
  freezing service groups 187
  freezing systems 221
  ignoreparent option 207
  linking resources 213
  linking service groups 194
  opening configuration files 222
  probing resources 208
  saving configuration files 222
  switching service groups 186
  taking resources offline 206
  taking resources offline and propagating 207
  taking service groups offline 186
  unfreezing service groups 188
  unfreezing systems 221
  unlinking resources 216
  unlinking service groups 195
commands, scripting 275
CompareRSM attribute 736
CompositeFileShare agent, configuring 298
ComputeStats attribute 703
conditional statements 258
ConfidenceLevel attribute 703
ConfigBlockCount attribute 728
ConfigCheckSum attribute 728
ConfigDiskState attribute 728
ConfigFile attribute 728
ConfigInfoCnt attribute 729
ConfigModDate attribute 729
configuration
  closing from Cluster Management Console 100

closing from Java Console 223
dumping 242
opening from Cluster Management Console 100
opening from Java Console 222
saving 242
saving from Cluster Management Console 100
saving from Java Console 222
saving in VCS Simulator 412
setting to read/write 242
setting to read-only 242
taking snapshots of 242
verifying 241
configuration files
generating 48
main.cf 48
read/write to read-only 250, 252, 253, 254, 255, 256, 261, 262
restoring from snaphots 242
taking snapshots of 242
types.cf 48
configuration language
local and global attributes 56
configurations
asymmetric 34
global cluster 45
N+1 38
N-to-1 36
N-to-N 40
replicated data 44
shared nothing 43
shared storage/replicated data 44
symmetric 35
configure
iSCSI initiator 278
ConfInterval attribute
about 459
definition 709
ConnectorState attribute 736
coordinator disks 429
CounterInterval attribute 737
CPU usage, how VCS monitors 632
CPUUsage attribute 729
CPUUsageMonitoring attribute 729
Critical attribute 703
CurrentCount attribute 717
CurrentLimits attribute 729
custom agents, about 26

D

Daemon Down Node Alive 438
DDNA 438
DeferAutoStart attribute 717
dependencies
between resources 126
for resources 22
for service groups
DHCP, disabling 283
Disk number, retrieving with DSRTest 678
DiskHbStatus attribute 729
domain analysis 95
domain, analyzing 95
dsrtest utility 677
dumping a configuration 242
DumpingMembership attribute 737
DynamicLoad attribute 730

E

ElifNone agent, configuring 341
Enabled attribute
for resources 703
for service groups 718
engine log
format 637
location 637
EnginePriority attribute 737
enterprise agents, about 26
entry points
about 25
modifying for performance 624
environment variables 58
error messages
agent log 637
at startup 654
engine log 637
message catalogs 639
Evacuate attribute 718
Evacuating attribute 718
event triggers
about 528
injeopardy 530
loadwarning 531
location of 528
nofailover 531
postoffline 532
postonline 532
preonline 533

resadminwait 533
resfault 535
resnotoff 536
resstatechange 537
sysoffline 538
unable_to_restart_had 538, 539
using 528
violation 539

F

failback, about 37
Failover attribute 718
FailOverPolicy attribute 718
FaultOnMonitorTimeouts attribute 710
file share groups
modifying using wizard 295
file shares, configuring using wizard 288
FileNone agent, configuring 341
FileOnOff agent, configuring 341
FileOnOnly agent, configuring 341
FileShare Configuration wizard 289
fire drills
about 565
for global clusters 565
for replicated data clusters 618
FireDrill attribute 710
Flags attribute 704
FromQ attribute 719
Frozen attribute
for service groups 719
for systems 730

G

GAB
about 422
impact on performance 622
verifying 651
when a system panics 629
GAB, about 27
gab_isolate_time timer 630
GenericService agent, configuring 326
getcomms utility 671
global attributes 56
global cluster configuration 45
global clusters
adding from Cluster Management Console 570
adding from Java Console 585
bringing remote groups online 594
deleting from Cluster Management Console 572
deleting from Java Console 589
operation 544
prerequisites for 551
setting up 553
switching remote groups 595
troubleshooting 667
user privileges 66
global heartbeats
administering from Cluster Management Console 578
administering from command line 614
administering from Java Console 596
deleting from Cluster Management Console 579
deleting from Java Console 599
modifying from Cluster Management Console 579
modifying from Java Console 598
global service groups
administering from Cluster Management Console 574
administering from command line 609
administering from Java Console 592
querying from command line 602
GlobalCounter attribute 738
Group Administrator
about 65
adding user as 249
Group attribute 704
group dependencies. See service group dependencies
Group Membership Services/Atomic Broadcast (GAB) 27
Group Operator
about 65
adding user as 250
GroupLimit attribute 738
GroupOwner attribute 719
Guests attribute
for clusters 738
for service groups 719
GUI. See Java Console or Cluster Management Console
GUIIPAddr attribute 730

H

haagent -display command 254

haagent -list command 258
haattr -add command 266
haattr -default command 267
haattr -delete command 267
hacf utility
  about 241
  creating multiple .cf files 241
  loading a configuration 241
  pretty-printing 241
hacf -verify command 241
HacliUserLevel attribute
  about 64
  definition 738
haclus -add command 612
haclus -declare command 612
haclus -delete command 612
haclus -display command
  for global clusters 606
  for local clusters 255
haclus -list command 606
haclus -modify command 612
haclus -state command 606
haclus -status command 606
haclus -value command
  for global clusters 606
  for local clusters 255
haclus -wait command 275
haconf -dump -makero command 242
haconf -makerw command 242
HAD
  about 27
  impact on performance 623
HAD Helper service, configuring manually 681
had -v command 273
had -version command 273
hadhelper command 681
hagetcf utility 672
hagrp -add command 259
hagrp -clear command 263
hagrp -delete command 259
hagrp -dep command 252
hagrp -disable command 262
hagrp -disableresources command 263
hagrp -display command
  for global clusters 602
  for local clusters 252
hagrp -enable command 262
hagrp -enableresources command 262
hagrp -freeze command 262
hagrp -list command
  for global clusters 603
  for local clusters 258
hagrp -modify command 259
hagrp -offline command
  for global clusters 610
  for local clusters 261
hagrp -online command
  for global clusters 609
  for local clusters 261
hagrp -resources command 252
hagrp -state command
  for global clusters 602
  for local clusters 252
hagrp -switch command
  for global clusters 610
  for local clusters 261
hagrp -unfreeze command 262
hagrp -unlink command 264
hagrp -value command 602
hagrp -wait command 275
hahb -add command 614
hahb -display command 607
hahb -list command 607
hamsg -info command 256
hamsg -list command 256
hanotify utility 516
hares -action command 611
hares -add command 264
hares -clear command 271
hares -delete command 266
hares -dep command 253
hares -display command
  for global clusters 604
  for local clusters 253
hares -global command 253
hares -info command 611
hares -link command 270
hares -list command
  for global clusters 604
  for local clusters 258
hares -local command 268
hares -modify command 266
hares -offline command 270
hares -offprop command 271
hares -online command 270
hares -override command 274
hares -probe command 271
hares -state command 604

hares -undo_override command 274
hares -unlink command 270
hares -value command 604
hares -wait command 275
hashadow process 27
hasnap -backup command 244
hasnap -delete command 248
hasnap -exclude command 247
hastart command 237
hastart -onenode command 237
hastart -ts command 238
hastatus command
    for global clusters 606
    for local clusters 255
hastatus -group command 255, 256
hastatus -summary command 256
hastop command 238
hasys -display command
    for global clusters 605
    for local clusters 255
hasys -freeze command 272
hasys -list command
    for global clusters 605
    for local clusters 255
hasys -modify command 272
hasys -nodeid command 272
hasys -state command 605
hasys -unfreeze command 272, 273
hasys -value command
    for global clusters 605
hasys -wait command 275
hatype -add command 274
hatype -delete command 274
hatype -display command 254
hatype -list command 254
hatype -modify command 274
hatype -resources command 254
hauser -add command 250
hauser -addpriv command 250
hauser -delete command 251
hauser -delpriv command 250, 251
hauser -display command 251
hauser -list command 251
havol utility 676
heap size for VRTSweb 766
Heartbeat 701
Heartbeat agent 547
heartbeat attributes 742
heartbeats, modifying for global clusters 614
host name issues 19

I

icons
    colors of 142
    in Java Console 138
IIS Configuration wizard 318
IIS configuration, synchronizing 317
IIS groups
    configuring using wizard 318
IIS sites, configuring 317
include clauses, about 49
InfoInterval attribute 710
InfoTimeout attribute 710
Initiator
    iSCSI 278
injeopardy event trigger 530
integer attribute type 53
IntentOnline attribute 720
IP agent, configuring 284
IPMultiNICPlus agent, configuring 285
iSCSI initiator
    configure 278
Istate attribute 705

J

Java Console
    about 30
    administering clusters 136
    administering logs 230
    administering resources 200
    administering service groups 178
    administering systems 219
    administering user profiles 173
    administering VCS Simulator 408
    arranging icons 155
    Cluster Explorer 146
    Cluster Manager 138
    Cluster Monitor 140
    Cluster Query 164
    components of 138
    customizing display 144
    icons 138
    impact on performance 625
    logging off of a cluster 172
    logging on to a cluster 170
    overview 136
    running commands from 223

- starting 137
- user profiles 173
- using with ssh 137
- viewing server credentials 167
- viewing user credentials 167

Java Console views
- Properties 151
- Remote Cluster Status 158
- Resource 153
- Service Group 152
- Status 150
- System Connectivity 157
- tear-off option 149

K

keylist attribute dimension 54
keywords 57
keywords, list of 57

L

Lanman agent, configuring 286
LastOnline attribute 705
LastSuccess attribute 720
LicenseType attribute 731
licensing issues 19
Limits attribute 731
LinkHbStatus attribute 731
LLT 28
- directives 648
- verifying 649

LLT, about 422
LLTNodeId attribute 731
Load attribute 720
Load policy for SGWM 452
LoadTimeCounter attribute 731
LoadTimeThreshold attribute 732
loadwarning event trigger 531
LoadWarningLevel attribute 732
local attributes 56
LockMemory attribute 738
LogDbg attribute 711
LogFileSize attribute 711
logging
- agent log 637
- engine log 637
- message tags 637
- VCW logs 638
- VCWsilent logs 638
- VRTSweb 761

logs
- customizing display in Java Console 230
- for VRTSweb 761
- searching from Java Console 230
- viewing from Java Console 165

LogSize attribute 739
Low Latency Transport (LLT) 28

M

main.cf
- about 48
- cluster definition 49
- group dependency clause 50
- include clauses 49
- resource definition 49
- resource dependency clause 49
- sample configuration 50
- service group definition 49
- system definition 49

ManageFaults attribute
- about 454
- definition 720

ManualOps attribute 720
message tags, about 637
Microsoft virtual machines, configuring 313
MigrateQ attribute 721
MonitorInterval attribute 711
MonitorOnly attribute 705
MonitorStartParam attribute 712
MonitorTimeout attribute 712
MonitorTimeStats attribute 706
MountV agent, configuring 281
MSMQ agent, configuring 332
MSMQ, configuring 332
MSVirtualMachine agent, configuring 313

N

N+1 configuration 38
Name attribute 706
network failure 157
network links, detecting failure 629
network resources, configuring 283
networks, detecting failure 630
NIC agent, configuring 284
NICTest utility 674
NoAutoDisable attribute 732
NodeId attribute 732

nodes
- adding to cluster 372
- removing from cluster 380

nofailover event trigger 531
notification
- about 512
- deleting messages 514
- error messages 513
- error severity levels 513
- event triggers 528
- hanotify utility 516
- message queue 514
- notifier process 515
- setting using wizard 227
- SNMP files 521
- troubleshooting 660

Notifier attribute 739
notifier process 515
Notifier Resource Configuration wizard 227
N-to-1 configuration 36
N-to-N configuration 40
NumRetries attribute 721
NumThreads attribute
- definition 712
- modifying for performance 624

O

OfflineMonitorInterval attribute 713
OfflineTimeout attribute 713
OnGrpCnt attribute 732
OnlineRetryInterval attribute 721
OnlineRetryLimit attribute
- for resource types 713
- for service groups 721

OnlineTimeout attribute 713
OnlineWaitLimit attribute 713
On-Off resource 22
On-Only resource 22
OpenTimeout attribute 713
Operations attribute 713
OperatorGroups attribute
- for clusters 739
- for service groups 721

Operators attribute
- for clusters 739
- for service groups 721

overload warning for SGWM 473

P

PanicOnNoMem attribute 739
Parallel attribute 722
passwords
- changing from Java Console 175

Path attribute 706
PathCount attribute 722
performance
- agents 623
- GAB 622
- HAD 623
- impact of VCS 622
- Java Console 625
- modifying entry points 624
- modifying NumThreads attribute 624
- monitoring CPU usage 632
- when a cluster is booted 626
- when a resource fails 628
- when a resource is brought online 626
- when a resource is taken offline 627
- when a service group fails over 631
- when a service group is brought online 627
- when a system fails 629

Persistent resource 22
Phantom agent, configuring 340
postoffline event trigger 532
postonline event trigger 532
PreOnline attribute 722
preonline event trigger 533
PreOnlineTimeout attribute 722
PreOnlining attribute 722
Prerequisites attribute 723
pretty-printing 241
print share groups
- creating using wizard 300
- modifying using wizard 311

print shares
- configuring using wizard 300

PrintMsg attribute 739
PrintShare Configuration wizard 300
PrintTree attribute 723
Priority attribute 723
privileges. See user privileges
Probed attribute
- for resources 706
- for service groups 723

ProbesPending attribute 723
Process agent, configuring 330
processes, configuring 330

ProcessPriority attribute 739
Proxy agent, configuring 340

Q

quick reopen 630

R

ReadOnly attribute 740
registry keys, excluding 337
registry replication, configuring 336
RegRep agent
- configuring 336
- keys for replication 337

Remote Cluster Configuration wizard 585
Remote Cluster States 694
remote clusters
- monitoring from Java Console 158
- states of 695

RemoteGroup agent, configuring 341
replicated data clusters
- about 44

replicated data configuration 44
resadminwait event trigger 533
reserved words 57
reserved words, list of 57
resfault event trigger 535
resnotoff event trigger 536
resource attributes 702
resource dependencies
- creating from Cluster Management Console 124
- creating from command line 270
- creating from Java Console 213
- displaying from command line 253
- removing from Cluster Management Console 125
- removing from command line 270
- removing from Java Console 215

resource faults
- clearing from Java Console 212
- simulating 412

resource type attributes 708
resource types
- importing 218
- querying from command line 254

ResourceInfo attribute
- clearing from Java Console 218
- definition 706
- refreshing from Java Console 218

ResourceLimit attribute 740
ResourceOwner attribute 707
resources
- about 21
- adding from Cluster Management Console 119
- adding from command line 264
- adding from Java Console 200
- administering from Cluster Management Console 118, 129
- administering from Java Console 200
- bringing online from Cluster Management Console 121
- bringing online from command line 270
- bringing online from Java Console 205
- categories of 22
- clearing faults from Cluster Management Console 123
- clearing faults from Java Console 212
- creating faults in VCS Simulator 412
- deleting from Cluster Management Console 124
- deleting from command line 266
- deleting from Java Console 205
- disabling from command line 468
- disabling from Java Console 211
- enabling from command line 262
- enabling from Java Console 210
- how disabling affects states 470
- invoking actions 217
- limitations of disabling 468
- linking from command line 270
- linking from Java Console 213
- On-Off 22
- On-Only 22
- Persistent 22
- probing from Cluster Management Console 124
- probing from Java Console 208
- querying from command line 253
- taking offline from command line 270
- taking offline from Java Console 206
- troubleshooting 659
- unlinking 126
- unlinking from Cluster Management Console 125
- unlinking from command line 270
- unlinking from Java Console 215

Responding attribute 723

resstatechange event trigger 537
Restart attribute 724
RestartLimit attribute
  about 459
  definition 714
root broker 29

S

saving a configuration 242
scalar attribute dimension 54
ScriptClass attribute 714
scripting VCS commands 275
ScriptPriority attribute 714
SCSI-III Persistent Reservations 431
search
  conducting 133
secure VCS. See VERITAS Security Services
Security Services
  configuring 78
seeding 645
server credentials, viewing 167
service group attributes 715
service group dependencies
  about 491
  autorestart 451
  benefits of 492
  categories of 493
  creating 506
  creating from Java Console 193
  limitations of 496
  manual switch 507
  removing from Java Console 195
service group workload management
  Capacity and Load attributes 472
  load policy 452
  load-based autostart 453
  overload warning 473
  sample configurations 475
  SystemZones attribute 452
service groups
  adding from Cluster Management Console 103, 105
  adding from command line 259
  adding from Java Console 178
  administering from Cluster Management Console 102
  administering from command line 259
  administering from Java Console 178
  autoenabling from Java Console 191
  bringing online from Cluster Management Console 108
  bringing online from command line 261
  bringing online from Java Console 183
  clearing from Cluster Management Console 115
  creating using configuration wizard 197
  deleting from Cluster Management Console 104
  deleting from command line 259
  deleting from Java Console 182
  disabling from Cluster Management Console 112
  disabling from Java Console 190
  displaying dependencies from command line 252
  enabling from Cluster Management Console 111
  enabling from Java Console 189
  flushing from Java Console 192
  freezing from command line 262
  freezing from Java Console 187
  linking from Java Console 193
  querying from command line 252
  switching from command line 261
  switching from Java Console 186
  taking offline from Java Console 185
  taking remote groups offline 595
  troubleshooting 656
  unfreezing from command line 262
  unfreezing from Java Console 188
  unlinking from Java Console 195
ServiceMonitor agent, configuring 327
services
  changing startup type 326
  configuring 326
shared nothing configuration 43
shared storage, configuring 278
shared storage/replicated data configuration 44
ShutdownTimeout attribute 733
Signaled attribute 707
Simulator. See VCS Simulator
SMTP notification, configuring for VRTSweb 755
SMTP server, retrieving name of 755
SNMP 512
  files for notification 521
  HP OpenView 521
  merging events with HP OpenView NNM 521
SNMP, supported consoles 512

SourceFile attribute
- for clusters 740
- for resource types 714
- for service groups 724
- for systems 733

split-brain
- in global clusters 549

ssh configuration for Java Console 137
Start attribute 707
State attribute
- for resources 707
- for service groups 724

steward process
- about 549
- configuring 559

Stewards attribute 740
Stopping 235
string attribute type 53
SupportedActions attribute 714
Symantec Product Authentication Service
- about 29
- authentication broker 29
- root broker 29
- viewing credentials 167

symmetric configuration 35
SysInfo attribute 733
SysName attribute 733
sysoffline event trigger 538
SysState attribute 733
System Attributes 728
system attributes 728
system states 696
SystemList attribute
- about 50, 260
- definition 724
- modifying 260

SystemLocation attribute 733
SystemOwner attribute 734
systems
- adding from Java Console 219
- adding to cluster 372
- administering from command line 272
- administering from Java Console 219
- bringing online in VCS Simulator 411
- client process failure 630
- deleting from Java Console 220
- detecting failure 629
- displaying node ID from command line 272
- freezing from Cluster Management Console 129
- freezing from Java Console 221
- panic 629
- quick reopen 630
- removing from cluster 380
- states 696
- unfreezing from Cluster Management Console 130
- unfreezing from Java Console 221

systems and nodes 20
SystemZones attribute 725

T

Tag attribute 725
TargetCount attribute 725
templates
- accessing Template View 159
- adding resources from 202
- adding service groups from 181

TFrozen attribute
- for service groups 725
- for systems 734

The 545
ToleranceLimit attribute 714
ToQ attribute 725
TriggerEvent attribute
- for resources 707
- for service groups 725

TriggerResFault attribute 726
TriggerResStateChange attribute 726
triggers. See event triggers
troubleshooting
- back up and restore files 666
- logging 637
- notification 660
- resources 659
- retrieving diagnostics 672
- service groups 656
- VCS startup 654
- Web Console 661

TRSE attribute 734
TypeDependencies attribute 726
TypeLimit attribute 740
types.cf 48

U

unable_to_restart_had trigger 538, 539
UpDownState attribute 734

user credentials, viewing 167
user privileges
- about 64
- assigning from command line 250
- changing from Java Console 176
- Cluster Administrator 64
- Cluster Guest 65
- Cluster Operator 65
- for specific commands 686
- Group Administrator 65
- Group Operator 65
- in global clusters 66
- removing from command line 250, 251

UserInt attribute 734
UserIntGlobal attribute 727
UserIntLocal attribute 727
UserNames attribute 740
users
- adding from Java Console 173
- deleting from command line 251
- deleting from Java Console 174
- displaying from command line 251

UserStrGlobal attribute 727
UserStrLocal attribute 727
utilities
- dsrtest 677
- getcomms 671
- hacf 241
- hagetcf 672
- hanotify 516
- havol 676
- NICTest 674
- VCSRegUtil 674
- vmgetdrive 680

V

VCS 568
- accessibility 767
- additional considerations for stopping 240
- assistive technology support 769
- event triggers 528
- logging 637
- notification 512
- querying from command line 252
- retrieving diagnostics with hagetcf 672
- seeding 645
- SNMP and SMTP 512
- starting as time-sharing process 238
- starting from command line 237
- starting on single node 237
- stopping from command line 238
- stopping with other options 239
- stopping without -force 239
- troubleshooting resources 659
- troubleshooting service groups 656
- using with VSFW 280
- verifying cluster operation 652
- verifying GAB 651
- verifying LLT 649

VCS agent statistics 633
VCS attributes 52
VCS Configuration wizard 73
VCS seeding 645
VCS Simulator
- administering from Java Console 408
- bringing systems online 411
- clearing cluster faults from Java Console 412
- creating power outages 412
- description of 276
- faulting resources 412
- installing 406
- saving offline configurations 412
- simulating cluster faults from command line 415
- simulating cluster faults from Java Console 410
- starting from command line 408

VCSFeatures attribute
- for clusters 741
- for systems 734

VCSi3Info attribute 740
VCSMode attribute 741
VCSRegUtil utility 674
VCW logs 638
VCWsilent logs 638
vector attribute dimension 54
Veritas 580
VERITAS Security Services
- disabling 391
- enabling 390

VERITAS Storage Foundation
- advanced features 280
- using with VCS 280

version information, retrieving 273
violation event trigger 539
virtual fire drill
- running 167

virtual names, configuring 286

VMDg agent, configuring 281
vmgetdrive utility 680
VRTSweb
adding ports 749
adding SMTP recipients 758
deleting ports 751
deleting SMTP recipients 760
logging 761
modifying log levels 762
notification for 755
ports for 748
retrieving log levels 761
retrieving ports 748
setting heap size 766
setting SMTP server 756
VRTSWebApp agent, configuring 339
vxfen. See fencing module

W

wac 546
WACPort attribute 741
Web Console
troubleshooting 661
wide-area connector 546
wide-area failover 45
Wide-Area Heartbeat agent 547
wizards
Application Configuration 343
FileShare Configuration 289
IIS Configuration 318
Notifier Resource Configuration 227
PrintShare Configuration 300
Remote Cluster Configuration 585
VCS Configuration 73