Veritas™ Cluster Server Release Notes

AIX 6.1

5.0 Maintenance Pack 1 Update 1

Veritas Cluster Server Release Notes 5.0 Maintenance Pack 1 Update 1 for AIX 6.1

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

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

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

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this product. Such third-party software is licensed separately by its copyright holder.

All third-party copyrights associated with this product are listed in the Veritas Cluster Server 5.0 Release Notes.

The Veritas Cluster Server 5.0 Release Notes can be viewed at the following URL:
http://support.veritas.com/docs/283978

Licensing and registration

Veritas Cluster Server is a licensed product. See the *Veritas Cluster Server Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit:
http://www.symantec.com/business/support/assistance_care.jsp.

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.

Veritas Cluster Server Release Notes

Introduction

This document provides important information regarding Veritas Cluster Server (VCS) version 5.0 MP1 Update 1 for AIX 6.1. Please review this entire document before installing or upgrading VCS.

For the latest information on updates, patches, and software issues regarding this release, see the following TechNote on the Veritas Technical Support website:

http://entsupport.symantec.com/docs/282024

Changes introduced in VCS 5.0 MP1 Update 1 for AIX 6.1

This section lists the changes in the 5.0 MP1 Update 1 for AIX 6.1 release of VCS.

Support for AIX 6.1

This release of VCS introduces support for AIX 6.1. Note that the VCS 5.0 MP1 patch was for previous versions of the AIX operating system.

See “System requirements” on page 21.

Table 1-1 describes the installation and upgrade options to move to AIX 6.1 with VCS 5.0 MP1 Update 1 for AIX 6.1.

Table 1-1 Available installation and upgrade paths

If you currently have...	To install or upgrade this MP, you need to... To move to AIX 6.1, you need to...
No installation of VCS	Use the disc to install VCS 5.0 MP1 Update 1 for AIX 6.1. See the *Veritas Cluster Server Installation Guide 5.0 for AIX*. See “Installing VCS 5.0 MP1 Update 1 for AIX 6.1” on page 24.
VCS 5.0 on AIX 5.3	Use the disc to upgrade to VCS 5.0 MP1 Update 1 for AIX 6.1. See “Upgrading to VCS 5.0 MP1 Update 1 for AIX 6.1” on page 24.
VCS 5.0 MP1 on AIX 5.3	Use the disc to upgrade to VCS 5.0 MP1 Update 1 for AIX 6.1. See “Upgrading to VCS 5.0 MP1 Update 1 for AIX 6.1” on page 24.

Continued support for AIX 5.3

This release of VCS continues support for AIX 5.3.

See "System requirements" on page 21.

Table 1-1 describes the installation and upgrade options for AIX 5.3 with VCS 5.0 MP1 Update 1 for AIX 6.1.

Table 1-2 Available installation and upgrade paths

If you currently have...	To install or upgrade this MP, you need to...
No installation of VCS	Use the disc to install VCS 5.0 MP1 Update 1 for AIX 6.1. See the *Veritas Cluster Server Installation Guide 5.0 for AIX*. See "Installing VCS 5.0 MP1 Update 1 for AIX 6.1" on page 24.
VCS 5.0 on AIX 5.3	Use the disc to upgrade to VCS 5.0 MP1 Update 1 for AIX 6.1. See "Upgrading to VCS 5.0 MP1 Update 1 for AIX 6.1" on page 24.
VCS 5.0 MP1 on AIX 5.3	Use the disc to upgrade to VCS 5.0 MP1 Update 1 for AIX 6.1. See "Upgrading to VCS 5.0 MP1 Update 1 for AIX 6.1" on page 24.

DNS agent supports A, AAAA records with corresponding PTR records

The DNS agent supports A and AAAA records. To accommodate this, and improve the agent, the HostName and Alias attributes have been replaced with the ResRecord attribute.

For details on the DNS agent:

See "DNS agent" on page 55.

Changes introduced in VCS 5.0 MP1

DNS agent supports secure DNS updates

The DNS agent has a new attribute called TSIGKeyFile. The attribute is required when you configure DNS for secure updates. The attribute specifies the absolute path to the file containing the private TSIG (Transaction Signature) key.

For details on the TSIG key:

See "Setting up secure updates using TSIG keys for BIND 9" on page 61.

Change in string size for some attribute values

For group name, resource name, attribute name, type name, and VCS username and password, the string size is limited to 1024 characters.

Updates to the NIC agent

For a virtual device, you must configure the NetworkHosts attribute. Symantec recommends configuring more than one host to take care of the NetworkHost itself failing. [781376]

Campus cluster support

You can configure a campus cluster using functionality provided by Veritas Volume Manager.

To set up a campus cluster, make sure the disk group contains mirrored volumes. The mirrors must be on separate storage at different sites. Use site tags to distinguish between mirrors located at different sites. You could also use enclosure-based naming. See the *Veritas Volume Manager Administrator's Guide* for detailed instructions.

Symantec recommends using I/O fencing in campus clusters.

Change in behavior: hastop command

VCS ignores the value of the cluster-level attribute EngineShutdown while the system is shutting down. [702597]

Change in behavior: BrokerIP attribute of the RemoteGroup agent

The BrokerIP attribute now requires only the IP address. Do not include the port number when you configure the attribute. [789878]

For a secure remote cluster only, if you need the RemoteGroup agent to communicate to a specific authentication broker, then set this attribute.

Type: string-scalar

Example: "128.11.245.51"

Fire drill support in Veritas Cluster Management Console

Veritas Cluster Management Console adds support for fire drills. The console lets you run fire drills and displays the status of the last fire drill.

- Viewing the status of the last fire drill—The service group listing tables display a column for the Physical Fire Drill Status, which indicates the results of the last fire drill.
- Running a fire drill.
  - Verify that replication for an application is working correctly
  - Verify that a secondary disaster recovery (DR) application service group can be brought online successfully.
- Viewing fire drill logs—If a service group is configured with a physical fire drill group, a tab labelled Fire Drill Logs appears on the secondary tab bar in the Group:Summary view. Click this tab to view the VCS log messages about the fire drill group on the remote cluster and the resources that belong to it.

See the *Veritas Cluster Server User's Guide* for information about fire drills.

Viewing the status of the last fire drill

The column Fire Drill Status has been added to service group listing tables. A service group listing table is on the Cluster:Groups view.

For VCS global service groups that are configured with a fire drill group, this column indicates the results of the most recently run fire drill. The following are the possible states:

UNKNOWN	No fire drill has been run or the Cluster Management Console has come online after the most recent fire drill
RUNNING	Fire drill in progress
PASSED	Fire drill group came online on the secondary cluster
FAILED	Fire drill group did not come online on the secondary cluster

If multiple management servers are connected to the global cluster that contains the primary global group, the table does not show fire drill status for that group.

Running a fire drill

The Cluster 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.

To run a fire drill from the Cluster 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.

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.

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.

Features introduced in VCS 5.0

See the *Veritas Cluster Server User's Guide* for details.

Cluster Management Console

The new Cluster Management Console replaces Cluster Manager (Web Console) and CommandCentral Availability.

Cluster Management Console enables administration and analysis for VCS clusters in your enterprise from a single console. You can install Cluster Management Console on a stand-alone system to manage multiple clusters or you can install the console on cluster nodes to manage a local cluster. When installed to manage a local cluster, the console is configured as part of the ClusterService group and the AppName attribute is set to `cmc`.

Cluster Monitor is now called Cluster Connector

CommandCentral Availability installed a component called Cluster Monitor on cluster nodes. The updated component is called Cluster Connector.

VCS privileges for operating system user groups

VCS 5.0 lets you assign VCS privileges to native users at an operating system (OS) user group level in secure clusters.

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.

See the *Veritas Cluster Server User's Guide* for more information.

Five levels of service group dependencies

VCS now supports configuring up to five levels of service group dependencies. The exception is the online local hard dependency, for which only two levels are supported.

New RemoteGroup agent to monitor service groups in remote clusters

The new RemoteGroup agent monitors and manages service groups in a remote cluster. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the agent.

Enhancements to the hastop command

You can customize the behavior of the hastop command by configuring the new EngineShutdown attribute for the cluster.

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; reject all other hastop commands.
PromptAlways	Prompt for user confirmation before running any hastop command.

Simulator supports deleting simulated clusters

VCS Simulator now supports deleting simulated clusters.

Symantec recommends using the same tool (command line or Java Console) to create and delete a cluster. For example, if you created the cluster from the Java Console, delete the cluster from the Java Console.

Fencing updates: DMP support

Dynamic multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. You can configure coordinator disks to use Veritas Volume Manager DMP feature.

You can set the coordinator disks to use either raw or DMP as the hardware path to a drive. See the *Veritas Cluster Server Installation Guide* for more information.

Minimal downtime upgrade to VCS 5.0

See the *Veritas Cluster Server Installation Guide* for a strategy on upgrading to VCS 5.0 while ensuring a minimal downtime for your applications.

Backup of VCS configuration files

VCS backs up all configuration files (<config>.cf) including main.cf and types.cf to <config>.cf.autobackup. The configuration is backed up only if the BackupInterval is set and the configuration is writable.

When you save a configuration, VCS saves the running configuration to the actual configuration file (i.e. <config>.cf) and removes all autobackup files. This does away with the VCS behavior of creating stale files

If you do not configure the BackupInterval attribute, VCS does not save the running configuration automatically.

See the *Veritas Cluster Server User's Guide* for more information.

Support for security services

VCS 5.0 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.

HAD diagnostics

When the VCS engine HAD dumps core, the core is written to the directory /var/VRTSvcs/diag/had, where the diagnostic information is stored. When HAD and GAB encounter heartbeat problems, VCS runs the script `/opt/VRTSvcs/bin/vcs_diag` to collect the diagnostic information.

The current working directory of VCS engine is VCS_DIAG whose default value is $VCS_HOME/diag. In earlier versions of VCS, the default directory of HAD was VCS_HOME whose default value was /opt/VRTSvcs.

Separate logger thread for HAD

The VCS engine, HAD, runs as a high-priority process to send heartbeats to kernel components and to respond quickly to failures. In VCS 5.0, HAD runs logging activities in a separate thread to reduce the performance impact on the engine due to logging.

Enhanced NFS lock failover

The new NFSRestart agent provides high availability to NFS locks. Use the agent in conjunction with the NFS agent. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.

Support for VLAN interfaces

The NIC and MultiNICA agents now support VLAN interfaces. The agents do not configure the NICs, but can monitor them.

See the OS vendor’s documentation on how to configure VLAN on your host, and ensure that the switch or router connected to such an interface is compatible with your configuration. Both server-side and switch-side VLAN configurations are supported.

Virtual fire drill

VCS supports a virtual fire drill capability that lets you test whether a resource can fail over to another node in the cluster. Virtual fire drills detect discrepancies between the VCS configuration and the underlying infrastructure on a node; discrepancies that might prevent a service group from going online on a specific node. See the *Veritas Cluster Server User’s Guide* for more information on running virtual fire drills.

New term: 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. See the *Veritas Cluster Server User’s Guide* for more information.

Change in behavior: Use comma or semicolon as delimiter

VCS 5.0 does not support using spaces as delimiters to separate vector, association, or keylist values. You must use a comma or a semicolon as a delimiter.

Change in behavior: New format for engine version

The new EngineVersion attribute replaces the MajorVersion and MinorVersion attributes. VCS stores version information in the following format:

```
<major>.<minor>.<maintenance_patch_num>.<point_patch_num>
```

Change in behavior for the resfault trigger

VCS now provides finer control over the resfault trigger. The resfault trigger is now invoked if the TriggerResFault attribute is set to 1.

Change in behavior: New location for enterprise agents

VCS enterprise agents are now installed in the /opt/VRTSagents/ha/bin directory.

The <agent>Types.cf files are now located at /etc/VRTSagents/ha/conf/<agent>.

Change in behavior: New location of message catalogs and attribute pools

VCS stores binary message catalogs (BMCs) at the following location:

/opt/VRTS/messages/*language*/module_name

The variable *language* represents a two-letter abbreviation.

The attribute pools also move from /var to /opt.

Change in behavior: New option for the hastart and had commands

Use the -v option to retrieve concise information about the VCS version. Use the -version option to get verbose information.

Changes to bundled agents

VCS introduces the following new agents:

- NFSRestart—Provides high availability for NFS record locks.
- RemoteGroup—Monitors and manages a service group on another system.
- Apache (now bundled on all platforms)—Provides high availability to an Apache Web server.

See “No longer supported” on page 23.

Changes to licensing for VCS

VCS now follows the licensing scheme that is described below:

License	What’s included
VCS	■ VCS ■ Cluster Management Console ■ Database agents ■ Application agents ■ Virtual fire drill support

License	What's included
VCS HA/DR	■ VCS ■ Cluster Management Console ■ Database agents ■ Application agents ■ Replication agents ■ Global clustering ■ Fire drill support

Note: Database agents are included on the VCS 5.0 disc. The replication and application agents are available via the Veritas High Availability Agent Pack.

New attributes

VCS 5.0 introduces the following new attributes. See the *Veritas Cluster Server User's Guide* for more information.

Resource type attributes

- AgentFile—Complete name and path of the binary for an agent. Use when the agent binaries are not installed at their default locations.
- AgentDirectory—Complete path of the directory in which the agent binary and scripts are located. Use when the agent binaries are not installed at their default locations.

Cluster attributes

- EngineShutdown—Provides finer control over the hastop command.
- BackupInterval—Time period in minutes after which VCS backs up configuration files.
- OperatorGroups—List of operating system user account groups that have Operator privileges on the cluster.
- AdministratorGroups—List of operating system user account groups that have administrative privileges on the cluster.
- Guests—List of users that have Guest privileges on the cluster.

System attributes

- EngineVersion—Specifies the major, minor, maintenance-patch, and point-patch version of VCS.

Service group attributes

- TriggerResFault—Defines whether VCS invokes the resfault trigger when a resource faults.
- AdministratorGroups—List of operating system user account groups that have administrative privileges on the service group.
- OperatorGroups—List of operating system user account groups that have Operator privileges on the service group.
- Guests—List of users that have Guest privileges on the service group.

Removed attributes

- DiskHbStatus—Deprecated. This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- MajorVersion—The EngineVersion attribute provides information about the VCS version.
- MinorVersion—The EngineVersion attribute provides information about the VCS version.

Updates to the DB2 agent

The Veritas High Availability Agent for DB2 introduces the following changes:

- The attributes StartUpOpt and ShutDownOpt provide new start up and shut down options. Using the StartUpOpt attribute, you can start the instance or partition, activate database commands after processes start, or create customized start up sequences. Using the ShutDownOpt attribute, you can perform a normal stop or customize your shut down sequence.
- In previous releases when you enabled in-depth monitoring (IndepthMonitor=1), it executed a default SQL query. The in-depth monitor now allows you to classify actions for DB2 errors according to their severity. You can associate predefined actions with each error code with a monitoring script that you can customize. You can find a sample of in-depth monitoring script in the following directory:
  /etc/VRTSagents/ha/conf/Db2udb/sample_db2udb.
  You must install the custom script in the /opt/VRTSagents/ha/bin/Db2udb directory to enable indepth monitoring.
- You can enable the AgentDebug attribute to get more debugging information from the agent and the database.

Updates to the Sybase agent

The Veritas High Availability Agent for Sybase agent supports Sybase ASE 12.5.x and 15 on AIX, HP-UX, Linux, and Solaris.

The agent supports encrypted passwords.

Updates to the Oracle agent

- New monitoring option—The basic monitoring option of the Oracle agent now allows health check monitoring in addition to the process check monitoring. You can choose the health check monitoring option for Oracle 10g and later.
- Support for virtual fire drills—VCS requires you to keep the configurations in sync with the underlying infrastructure on a cluster node. Virtual fire drills detect such discrepancies that prevent a service group from going online on a specific system. Refer to the *Veritas Cluster Server User's Guide* for more information.
  The agent uses the Action entry point to support the virtual fire drill functionality.

About Veritas agents

Veritas Cluster Server bundled agents

VCS bundles agents to manage a cluster's key resources. The implementation and configuration of bundled agents vary by platform.

For more information about bundled agents, refer to the *Veritas Cluster Server Bundled Agent Reference Guide*.

Veritas High Availability Agent Pack

The Veritas High Availability Agent Pack gives you access to agents that provide high availability for third-party storage solutions. It is re-released regularly to add new agents. Contact your Symantec sales representative for information about agents included in the agent pack, agents under development, and agents that are available through Symantec consulting services.

Veritas Cluster Server custom agents

VCS provides a framework that allows for the creation of custom agents. Create agents in situations where the Veritas High Availability Agent Pack, the bundled agents, or the agents for enterprise applications do not meet your needs. You can also request a custom agent through Symantec consulting services.

For more information about the creation of custom agents, refer to the *Veritas Cluster Server Agent Developer's Guide*.

Veritas Cluster Server agents for enterprise applications

VCS also provides agents to manage key enterprise applications. This section lists the agents for enterprise applications and the software that the agents support.

Note: Before configuring an enterprise agent with VCS, verify that you have a supported version of the agent.

Veritas agents support a specified application version on AIX if the application vendor supports that version on AIX.

Agent	Agent version	VCS version			Application		OS	
		3.5	4.0	5.0			5.3	6.1
DB2	5.0	p	p	s	DB2 Enterprise Server Edition	8.1, 8.2, 9.1	s	s
Oracle	5.0	p	p	s	Oracle	9*i*, 10g R1, 10g R2	s	s
Sybase	5.0	p	p	s	Sybase Adaptive Server Enterprise	12.5.x, 15	s	s

s – supported configuration p – supported by previous version of the agent

Custom agents

Custom agents developed in C++ must be compiled using the IBM C for AIX Compiler Version 7.0. Use the -`brtl` flag for runtime linking with the framework library.

System requirements

This section describes system requirements for VCS.

Before you install your Symantec products, you must read the Late Breaking News TechNote:

http://entsupport.symantec.com/docs/282024

Supported hardware

The compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:

http://entsupport.symantec.com/docs/283161

Before installing or upgrading Veritas Cluster Server, review the current compatibility list to confirm the compatibility of your hardware and software.

Required patch

Before you install your Symantec products, you must read the following TechNote and perform the instructions in it.

http://entsupport.symantec.com/docs/300577

Supported software

- AIX 5.3 TL7 with SP2 is required
  See "Required patch" on page 21.
- AIX 6.1 with SP3 and the required iFix
  See "Required patch" on page 21.
- Logical Volume Manager (LVM)
- Journaled File System (JFS) and Enhanced Journaled File System (JFS2)
- Veritas Volume Manager (VxVM) 5.0 MP1 Update 1 for AIX 6.1
- Veritas File System (VxFS) 5.0 MP1 Update 1 for AIX 6.1

Supported software for Cluster Management Console

You can install Cluster Management Console on a standalone system to manage multiple clusters or you can install the console on cluster nodes to manage a local cluster.

Supported browsers

Veritas Cluster Management Console is supported on the following browsers:

- Microsoft Internet Explorer 6.0 with SP2 or later
- Firefox 1.5 or newer

Veritas Cluster Management requires the Macromedia Flash Plugin v8.0.

Requirements for accessing Cluster Manager (Java Console)

Cluster Manager (Java Console)

The VCS Java Console requires a minimum of 256MB RAM and 1280x1024 display resolution. The color depth of the monitor must be at least 8-bit (256 colors), although 24-bit is recommended.

The minimum requirements for Windows clients are Pentium II, 300MHz, 256MB RAM, and 800x600 display resolution. (Symantec recommends a minimum of Pentium III, 400MHz, and 512MB RAM.) The color depth of the monitor must be at least 8-bit (256 colors), and the graphics card must be able to render 2D images.

Cluster Manager requires AIX developer kit for Java

Cluster Manager (Web Console and Java Console) requires the IBM AIX Developer Kit, Java 2 Technology Edition, Version 1.3.0.

No longer supported

Support is no longer provided for:

- ServiceGroupHB agent. This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- Disk heartbeats (GABDisk). This release does not support disk heartbeats. Symantec recommends using I/O fencing.
- The updated Oracle agent does not support Oracle 8.0.x and Oracle 8.1.x.
- The updated DB2 Agent does not support DB2 7.2

Installation notes

Installing VCS 5.0 MP1 Update 1 for AIX 6.1

With this release's maintenance pack disc, perform the instructions in the *Veritas Cluster Server Installation Guide 5.0 for AIX* for a new installation of VCS and this maintenance pack. Review and perform the instructions in the following chapters:

- Preparing to install and configure VCS
- Installing and configuring VCS

Upgrading to VCS 5.0 MP1 Update 1 for AIX 6.1

To upgrade from VCS 5.0 or VCS 5.0 MP1, perform the following tasks:

- Preparing to upgrade
- Running the upgrade program
- Performing post-upgrade tasks
- Migrating the operating system from AIX 5.3 to AIX 6.1

Preparing to upgrade

1. Log in as superuser on one of the nodes in the cluster.
2. Insert the disc that contains the 5.0 MP1 Update 1 for AIX 6.1 software into the disc drive of one of the nodes in the cluster.
3. Mount the disc on a suitable mount point.
4. Navigate to the highest-level directory that contains the installmp program.

5 Verify that /opt/VRTSvcs/bin is in your PATH. Do one of the following:

- For the Bourne Shell (sh or ksh), type:

  ```
  $ PATH=/usr/sbin:/sbin:/opt/VRTS/bin:/opt/VRTSvcs/bin:\
  $PATH; export PATH
  ```

- For the C Shell (csh or tcsh), type:

  ```
  % setenv PATH /usr/sbin:/sbin:/opt/VRTS/bin:\
  /opt/VRTSvcs/bin:$PATH
  ```

6 If VCS is running and service groups are configured, enter the `hastop -all` command to stop VCS:

```
# hastop -all
```

Do not use the `-force` option with the `hastop` command. Using the force option leaves all service groups online and can cause undesired results after the upgrade.

Running the upgrade program

1 Install 5.0 MP1 Update 1 for AIX 6.1 using the installmp program:

```
./installmp [-rsh]
```

2 After the initial system checks and the requirements checks are complete, press the Return key to start upgrading the filesets.

3 When the installation completes, note the locations of the summary, log, and response files.

Performing post-upgrade tasks

1 Update the types.cf file to the new version.

```
cp -p /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.orig
cp -p /etc/VRTSvcs/conf/types.cf \
/etc/VRTSvcs/conf/config/types.cf
```

2 If you had added custom type definitions in the original types.cf file, add them to the new types.cf file.

3 If you are migrating from AIX 5.3 to AIX 6.1, do not start any VCS processes at this time. Skip to:

See “Migrating the operating system from AIX 5.3 to AIX 6.1” on page 26.

4 If you are upgrading on AIX 5.3:

- Update the main.cf file to configure resources affected by the upgrade. See “Changes introduced in VCS 5.0 MP1 Update 1 for AIX 6.1” on page 6.

- Execute the following command to restart your systems:

  ```
  /usr/sbin/shutdown -r
  ```

Migrating the operating system from AIX 5.3 to AIX 6.1

1. If you want to migrate the operating system to AIX 6.1, perform the operating system migration from AIX 5.3 to AIX 6.1.
2. After you have completed a successful operating system migration, the VCS nodes reboot.

   GAB, LLT and VCS processes start automatically and previously configured service groups come online.

Installation notes for VCS 5.0

The following information includes guidelines, tips, and other considerations for installing VCS 5.0.

For more information, refer to the *Veritas Cluster Server Installation Guide*.

Before upgrading VCS to version 5.0

If you plan to upgrade VCS with NFS configuration from versions earlier than 5.0, you must perform the following pre-upgrade tasks:

1. Take a backup of the main.cf file.
2. Remove the NFS configuration from the main.cf.
3. Upgrade VCS to 5.0.
4. After you upgrade VCS to 5.0, add the NFS configuration in the main.cf file.

Change default password after installing VCS

When you install and configure VCS, if you do not choose the secure mode, the installvcs program creates a user *admin* with the password *password*. The user has administrative privileges to the cluster.

Symantec recommends you change the password of the user after installing and configuring VCS. See the *Veritas Cluster Server User's Guide* for more information.

If you used the AllowNativeCliUsers attribute

If you used the AllowNativeCliUsers attribute, see the *Veritas Cluster Server Installation Guide* for information on how to use the halogin utility after upgrading to VCS 5.0.

Upgrading the VCS Java Console

This release includes updates for Cluster Manager (Java Console)

To upgrade the Java Console on a Windows client

1. Stop Cluster Manager if it is running.
2. Remove Cluster Manager from the system.
3. Insert the software disc into a drive on your Windows system.
4. Start the installer from the following path:
   \windows\VCSWindowsInstallers\ClusterManager\EN\setup.exe
5. Follow the wizard instructions to complete the installation.

Upgrading the VCS Simulator

This release includes updates for VCS Simulator.

To upgrade VCS Simulator on a Windows client

1. Stop all instances of VCS Simulator.
2. Stop VCS Simulator, if it is running.
3. Remove VCS Simulator from the system.
4. Insert the software disc into a drive on your Windows system.
5. Start the installer from the following path:
   \windows\VCSWindowsInstallers\Simulator \EN\vrtsvcssim.msi
6. Follow the wizard instructions to complete the installation.

Installing 5.0 MP1 Update 1 for AIX 6.1 VCS Agent for Sybase

If you are installing the VCS Agent for Sybase on a cluster that has no previous installation of the Sybase agent, perform the following procedure.

To install the VCS AGent for Sybase for 5.0 MP1 Update 1 for AIX 6.1

1. On the mounted disc, change directory to the cluster_server_agents/sybase_agent/pkgs.
2. Install the Sybase fileset. This directory contains the 5.0.0.0 version.
3. Change directory to cluster_server_agents/sybase_agent/patches.
4. Install the Sybase fileset. This directory contains the 5.0.1.100 version.

Removing VCS 5.0 MP1 Update 1 for AIX 6.1

VCS 5.0 MP1 Update 1 for AIX 6.1 cannot be uninstalled separately from VCS. You must completely uninstall VCS if you no longer want the patch. Use the installation program to uninstall VCS.

To uninstall VCS

1. Confirm that you are logged in as the superuser and mounted the product disc.
2. Start the installer.

   ```
   # ./installer
   ```
3. The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.
4. From the opening Selection Menu, choose: “u” for “Uninstall a Product.”
5. From the displayed list of products to install, choose the number that corresponds to Veritas Cluster Server.
6. Reply to the prompts to uninstall the product.

Fixed issues

Refer to the following sections for information about:

- Issues fixed in VCS 5.0 MP1 Update 1 for AIX 6.1
- Issues fixed in VCS 5.0 MP1
- Issues fixed in VCS 5.0

Issues fixed in VCS 5.0 MP1 Update 1 for AIX 6.1

The following issues were fixed in this release. For a list of additional issues fixed in this release, see the following TechNote:
http://entsupport.symantec.com/docs/285869

1218881	The Apache version was incorrectly parsed for IBMIHS.
1206153	Updated DNS name validation to better conform to RFC definition.
1201174	To resolve an inconsistent HAD core dump in VCSMutexDestroy, HAD now always calls VCSExit not exit.
1199313	Increased the binary capacity to accommodate larger configurations (MAXDATA increase for HAD).
1187827	To resolve online and monitor issues, changed DNS agent.
1187580	Resolved certain ActionTimeout attribute issues.
1186414	Repaired a problem where hastart and triggers needed to run on the locale specified by the LANG variable.
1180976	Fixed esballoc and allocb calls to not modify internal flags that can cause a panic under stress conditions.
1177984	Resolved an issue where the `haclus -modify BackupInterval` hung when resetting the value to the same value.
1174911	Group switch/failover logic now completes if parent group gets autodisabled during fail over.
1174520	If a ServiceGroup with a single system in a SystemList is unfrozen, HAD no longer asserts.
1168476	Multiple MultiNICA issues resolved (@netHosts array, getAliases subroutine, and %niclist hash).
1161339	The Application agent now inherits the user defined LANG parameter.
1139831	Support SyncODM attribute for scalable volume groups added.

1116813 In the Java Console, fixed an issue where the Global group wizard gives lowest priority to the local (primary in some cases) cluster.

1116600 The Oracle Agent does not correctly identify state if the SID value is same for two resources in two different containers.

1113667 Replaced a fork call with a safer one.

1105310 When msgid 13319 is added using the `halog -add` command, the core no longer dumps.

1096394 Fixed an issue where notifier blocked the IPM socket.

1093791 Fixed an issue where it took a long time to detect a secondary site fault after the loss of wac and HAD.

1056559 Fixed an issue where the NFSRestart monitor threw a, “too many open files error.”

1053377 Changed the path value so that VCS picks up the correct df etc values.

1050999 Resolved a HAD issue that caused the ShutdownTimeout value to not work correctly.

1016548 Fixed an issue where a node panicked with a GAB: Port f halting system due to network failure message.

1016532 Fixed an issue where pings saw a 50% packet loss after a successful local failover when an interface goes down.

1014281 Fixed an issue where Oracle encountered vcsipc send error 12, not enough space, under high interconnect traffic.

995504 Fixed an issue where haping intermittently timed out with an error.

995199 The Process agent no longer dumps core if more than two Process resources are configured.

862507 When the sequence request is unsent, GAB_F_SEQBUSY is not set.

313089 The MultiNICB agent now moves the default route to an active interface after a failover.

Issues fixed in VCS 5.0 MP1

The following issues were fixed in this release. For a list of additional issues fixed in this release, see the following TechNote:
http://entsupport.symantec.com/docs/285869

805121	Partial groups go online erroneously if you kill and restart the VCS engine.
796655	The DNS agent does not support secure updates.
784335	The Oracle agent cannot identify the shell when the /etc/passwd file has multiple occurrence of the $Owner string.
781376	The NIC agent fails to work in virtual ethernet environments.
774893	The fencing module panics if scsi_disk_policy is set to dmp.
702597	VCS ignores the value of the cluster-level attribute EngineShutdown while the system is shutting down.
702594	The Oracle agent does export SHLIB_PATH and other environment in CSH.
646372	The `hatype -modify` ... `-delete` ...command works incorrectly. The command deletes the first element of the keylist attribute.
627647	The Action entry point for Oracle fails because set_environment() function prototype differs.
627568	The STARTUP_FORCE value needs to be added in the drop-down list of StartUpOpt values in the Oracle and RAC wizards as the default value for StartUpOpt.
625490	For the agent framework module, ag_i18n_inc.sh does not invoke halog when script entry points use the VCSAG_LOGDBG_MSG API, even if the debug tag is enabled.
620529	Cluster Management Console does not display localized logs. If you installed language packs on the management server and on VCS 5.0 cluster nodes, Cluster Management Console did not initially show localized logs.
619219	Running the hastart command twice causes an assertion to be displayed.
616964	In a secure environment, the RemoteGroup agent does not authenticate on a remote host for the first time.
616580	Importing resource types fails on Simulator on Windows systems.

609555	The Remote Group Agent wizard in the Java GUI rejects the connection information for the remote cluster with the domain type other than the local cluster. Fix: The RGA Wizard can now connect to all supported domain types irrespective of the domain type of local cluster.
608926	The template file for the DB2 agent does not contain the complete information for building a DB2 MPP configuration. The template does not include a service group required in the configuration.
598476	If you have a service group with the name ClusterService online on the last running node on the cluster, the hasim -stop command appears to hang.
570992	Cluster Management Console does not display some icons properly.
565151	The NIC agent does not detect failure when you pull the network cable.
545469	The Monitor entry point does not detect an online when the Oracle instance is not started by the user defined in the Owner attribute.
244988	Very large login name and password takes all the service groups offline. Fix: For group name, resource name, attribute name, type name, and VCS username and password, the string size is limited to 1024 characters.
243186	Assertion in VCS engine.

Issues fixed in VCS 5.0

The following issues were fixed in VCS 5.0.

	The concurrency violation trigger could not offline a service group if the group had a parent online on the system with local firm dependency. The concurrency violation continued until the parent was manually taken offline.
	The configuration page for the Symantec Web server (VRTSWeb) offered two Japanese locale options. Both options had UTF-8 encoding, and there were no functional difference between the two.
	The agent for Oracle obtained its initialization parameters from the pfile. VCS could not monitor Oracle instances created from the spfile.
	When installing Cluster Manager on a Windows XP system, the following error appeared: “The installer has insufficient privileges to access this directory: C:\Config.Msi.”
648584	Made the LLT timer and service threads’ priorities a tunable entity.

620378	Complex group dependencies and timing issues leads to different failovers.
584243	hares options do not filter correctly.
520034	Issues supporting multiple instances of MultiNICA agent.
515644	hacf does not handle MAXARG values of vector/associative attributes in the main.cf.
426932	Indeterministic service thread cancellation.
418971	Cannot configure multiple Sybase servers with VCS.
393849	Performance issues with the Mount agent.
297779	Support multiple MultiNICB instances.
271167	Provide finer control over the hastop -all command.
254947	GAB and LLT device files have open permissions.
252347	Behavior of parent group is incorrect when groups are linked with online global firm and child group faults.
248069	Commands do not close socket after successful termination.
247698	Need to move logging activities out of single-threaded HAD.
246238	Information required when had is restarted either by hashadow or gab.

Known issues

The following issues are open for this version of VCS.

Operational issues for VCS

Question marks in Application agent logs

Questions marks may appear in the Application agent logs. You can safely ignore them. [1220322]

Saving large configuration results in very large file size for main.cf

If your service groups have a large number resources or resource dependencies, and if the PrintTree attribute is set to 1, saving the configuration may cause cause the configuration file to become excessively large in size and may impact performance. [616818]

Workaround: Disable printing of resource trees in regenerated configuration files by setting the PrintTree attribute to 0.

AutoStart may violate limits and prerequisites load policy

The load failover policy of Service Group Workload Management may be violated during AutoStart when all of the following conditions are met:

- More than one autostart group uses the same Prerequisites.
- One group, G2, is already online on a node outside of VCS control, and the other group, G1, is offline when VCS is started on the node.
- The offline group is probed before the online group is probed.

In this scenario, VCS may choose the node where group G2 is online as the AutoStart node for group G1 even though the Prerequisites load policy for group G1 is not satisfied on that node.

Workaround: Persistently freeze all groups that share the same Prerequisites before using `hastop -force` to stop the cluster or node where any such group is online. This workaround is not required if the cluster or node is stopped without the force option.

Trigger not invoked in REMOTE_BUILD state

In some situations, VCS does not invoke the injeopardy trigger if the system is a REMOTE_BUILD state. VCS fires the trigger when the system goes to the RUNNING state.

Issue with offline local group dependencies

This issue occurs with offline local group dependencies, when the parent service group has the AutoFailover attribute set to 0. When the child group faults, it does not fail over to system where parent is online, even though that is the only system available for failover. [248532]

The hacf -verify command fails if the IP address is specified

The hacf -verify command fails if you specify the IP address instead of the system name. [834496]

Node cannot join cluster because port v is not ready for configuration

This behavior is observed when a node leaves a cluster and another node tries to join the cluster at the same time. If the GAB thread is stuck in another process, the new node cannot join the cluster and GAB logs the following warning:

```
GAB WARNING V-15-1-20126 Port v not ready for reconfiguration, will
retry.
```

The haclus -wait command hangs when cluster name is not specified

If you do not specify the cluster name when running the `haclus -wait` command, the haclus -wait command may hang. [612587]

Using the coordinator attribute

This release contains an attribute for disk groups called coordinator, which configures disks as coordinator disks by the I/O fencing driver. Setting the attribute prevents the coordinator disks from being reassigned to other disk groups. See the Veritas Volume Manager documentation for additional information about the coordinator attribute.

The attribute requires that the disk group contain an odd number of disks. Symantec recommends that you use only three coordinator disks. Using more (five or seven) disks may result in different subclusters.

HAD cannot join clusters if ulimit is low

The default ulimit may restrict the amount of memory that can be allocated by a process. Run the following command to verify this limit:

```
ulimit -a
  data seg size   (kbytes -d) 131072
```

In case of very large main.cf configurations (example 10,000 lines) or multi-node large clusters, or a combination of these, the VCS engine HAD fails to join clusters. This could be a result of HAD not being able to allocate enough memory to facilitate transfer of main.cf across cluster nodes.

Workaround: Increase the data seg size to an appropriate value by running the `ulimit -d` *`value`* command.

Some alert messages do not display correctly

The following alert messages do not display correctly [612268]:

51030	Unable to find a suitable remote failover target for global group %s. administrative action is require
51031	Unable to automatically fail over global group %s remotely because local cluster does not have Authority for the group.
50913	Unable to automatically fail over global group %s remotely because clusters are disconnected and ClusterFailOverPolicy is set to %s. Administrative action is required.
50914	Global group %s is unable to failover within cluster %s and ClusterFailOverPolicy is set to %s. Administrative action is required.
50916	Unable to automatically failover global group %s remotely due to inability to communicate with remote clusters. Please check WAN connection and state of wide area connector.
50761	Unable to automatically fail over global group %s remotely because ClusterList values for the group differ between the clusters. Administrative action is required.
50836	Remote cluster %s has faulted. Administrative action is required.
51032	Parallel global group %s faulted on system %s and is unable to failover within cluster %s. However, group is still online/partial on one or more systems in the cluster
51033	Global group %s is unable to failover within cluster %s and AutoFailOver is %s. Administrative action is required.

Issues related to the VCS engine

GAB may kill HAD

In certain situations of heavy load on nodes, GAB can kill HAD. [1200256]

Engine may hang in LEAVING state

When the command `hares -online` is issued for a parent resource when a child resource faults, and the `hares -online` command is followed by the command `hastop -local` on the same node, then the engine transitions to the LEAVING state and hangs.
Workaround: Issue the command `hastop -local -force`.

Timing issues with AutoStart policy

Consider a case where the service group is offline and engine is not running on node 1. If you restart the engine on node 1 after HAD is killed on node 2 *and* before the engine is restarted on node 2, then VCS does not initiate the autostart policy of the group.

Issues related to fencing

Preexisting split brain after rebooting nodes

The fencing driver in 5.0 uses Veritas DMP to handle SCSI commands to the disk driver if fencing is configured in dmp mode. This allows fencing to use Veritas DMP for access to the coordinator disks. With certain disk arrays, when paths are failed over due to a path failure, the SCSI-3 persistent reservation keys for the previously active paths are not removed. If the nodes in a cluster are all rebooted at the same time, then the cluster will not start due to a `Preexisting split brain` message. [609407]

Workaround: Use the `vxfenclearpre` script to remove the keys from the coordinator disks as well as from the data disks.

Stopping vxfen when the fencing module is being configured

Trying to stop the vxfen driver when the fencing module is being configured results in the following error.

```
VCS FEN vxfenconfig ERROR V-11-2-1013 Unable to unconfigure vxfen
VCS FEN vxfenconfig ERROR V-11-2-1022 Active cluster is currently
fencing.
```

Workaround: This message may be safely ignored.

Fencing configuration fails if fencing module is running on another node

The `vxfenconfig -c` command fails if any of the following commands are running on other nodes in the cluster:

```
vxfenconfig -U
vxfenconfig -c
```

Some vxfenadm options do not work with DMP paths

Some options of the vxfenadm utility do not work well with DMP paths such as /dev/vx/rdmp/sdt3.

Workaround: Use the -a option to register keys instead of -m option for DMP paths.

Issues related to global service groups

Switch across clusters may cause concurrency violation

If you try to switch a global group across clusters while the group is in the process of switching across systems within the local cluster, then the group may go online on both the local and remote clusters. This issue affects only global groups. Local groups do not experience this behavior.

Workaround: Ensure that the group is not switching locally before attempting to switch the group remotely.

Global service group does not go online on AutoStart node

At cluster startup, if the last system where the global group is probed is not part of the group's AutoStartList, then the group does not AutoStart in the cluster. This issue affects only global groups. Local groups do not display this behavior.

Workaround: Ensure that the last system to join the cluster is a system in the group's AutoStartList.

Declare cluster dialog may not display highest priority cluster as failover target

When a global cluster fault occurs, the Declare Cluster dialog enables you to fail groups over to the local cluster. However, the local cluster may not be the cluster assigned highest priority in the cluster list.
Workaround: To bring a global group online on a remote cluster, do one of the following:

- From the Java Console, right-click the global group in the Cluster Explorer tree or Service Group View, and use the Remote Online operation to bring the group online on a remote cluster.
- From the Web Console, use the Operations links available on the Service Groups page to bring the global group online on a remote cluster.

The gcoconfig command assigns priority 0 to all nodes

If you configure a global cluster using the `/opt/VRTSvcs/bin/gcoconfig` command, the gcoconfig utility assigns the same priority '0' to all the nodes that are in the SystemList of the ClusterService group. [857159]

Workaround: Edit main.cf and assign priority for cluster nodes in the SystemList of the ClusterService group.

Use one of the following approaches to edit the main.cf file:

- Veritas Cluster Server GUI
- VCS commands
- Stop VCS and manually edit the main.cf file
  Note that this approach has HA downtime.

Issues related to VCS bundled agents

Problem in failing over the IP resource

When a system panics, the IP address remains plumbed to the system for a while. In such a case, VCS may not succeed in failing over the IP resource to another system. This can be observed when a system panics during I/O Fencing.

Workaround: Increase the value of the OnlineRetryLimit attribute for the IP resource type.

Taking a group with the Mount resource offline can take several minutes if the file system is busy

When a file system has heavy I/O, the umount command can take several minutes to respond. However, the umount command deletes the mount point from mount command output before returning. Per IBM, this is the expected and supported behavior on AIX. The umount command's processing later puts the mount point back if the mount point is found busy. Meanwhile, the default OfflineTimeout value of the Mount agent can get exceeded, which in turn invokes the Clean agent function. The Clean function can find the mount point's entry absent from the mount command output and exit with success.

The unmounting, however, may not have happened yet. If unmounting did not occur, offlining resources below the Mount resource (for example the LVMVG or DiskGroup resources) can fail.

The Mount resource's Offline agent function then proceeds to unmount the mount point. After several attempts, the Clean scripts that clean the resources below the Mount resource succeed and the group goes offline.

See the *Veritas Cluster Server User's Guide* for more information about the OfflineTimeout attribute.

LVMVG agent with big and scalable volume groups

For big and scalable volume groups, the LVMVG agent does not properly synchronize the ODM.

Workaround: Set the attribute `SyncODM = 0` and manually synchronize the ODM when adding a volume group.

Issues related to the DB2 agent

All partitions fault even if there are errors on only one partition with the IndepthMonitor database

This issue occurs in an MPP environment when multiple partitions use the same database. If the Databasename attribute is changed to an incorrect value, all partitions using the database fault. [568887]

Db2udb resource faults when IndepthMonitor is configured with a Japanese database

For locales other than English, you need to add the following lines to the $INSTHOME/sqllib/userprofile file. [590010]

The following example adds Japanese support for AIX:

```
export LANG=Ja_JP
```

Issues related to the Oracle agent

NOFAILOVER action specified for certain Oracle errors

The Veritas High Availability agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file oraerror.dat, which consists of a list of Oracle errors and the actions to be taken. Refer to the *Veritas High Availability Agent for Oracle Installation and Configuration Guide* for a description of the actions.

Currently, the reference file specifies the NOFAILOVER action when the following Oracle errors are encountered:

ORA-00061, ORA-02726, ORA-6108, ORA-06114

The NOFAILOVER action means that the agent sets the resource's state to OFFLINE and freezes the service group. You may stop the agent, edit the oraerror.dat file, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

Health check may not work

If you set MonitorOption to 1, health check monitoring may not function when the following message is displayed [589934]:

```
Warning message - Output after executing Oracle Health Check is:
GIM-00105: Shared memory region is corrupted.
```

Workaround: Set MonitorOption to 0 to continue monitoring the resource.

Health check monitoring does not work in csh environment

If you use csh, you must change your shell environment to enable the Health check monitoring option.

Issues related to Cluster Manager (Java Console)

Exception when selecting preferences

On Windows systems, selecting the Java (Metal) look and feel of the Java Console may cause a Java exception. [585532]

Workaround: After customizing the look and feel, close restart the Java Console.

Java Console errors in a localized environment

When connected to cluster systems using locales other than English, the Java Console does not allow importing resource types or loading templates from localized directories. [585532]

Workaround: The workaround is to copy the types files or templates to directories with english names and then perform the operation.

Printing to file from the VCS Java Console throws exception

VCS Java Console and Help throw an exception while printing to a file from a system that does not have a printer configured. Also, the content is not written to the file.

Workaround: Before printing, make sure at least one printer is configured on the system where the VCS Java Console is launched.

Common system names in a global cluster setup

If both local and remote systems have a common system name in a global cluster setup, group operations cannot be performed on those systems using the Java console.

Workaround: Use command-line interface to perform group operations.

Issues related to Cluster Management Console

Installation using a response file for a cluster enabled with Symantec Product Authentication Service does not create the CMC service group

Installing VCS using a response file (using the installer program with the -responsefile option) for a cluster that uses Symantec Product Authentication Service does not create the CMC service group.

Workaround: Perform the installation with the response file. After installation, manually edit the main.cf file to add CMC-related types and resources. You can then bring theses service groups online. [1218568]

Wrong file size error for ClusterConnectorConfigType.cf file

f you use the 5.0 MP1 installmp script to apply the maintenance patch to the VRTScmccc.rte file set, and then reject the patches using the installp -r VRTScmccc.rte command, you get following error messages: [806706]

```
sysck: 3001-049 Wrong file size. The file
/etc/VRTSvcs/conf/config/ClusterConnectorConfigType.cf has an
actual size of 754 bytes (expected size: 789 bytes).
sysck: 3001-017 Errors were detected validating the files for
package VRTScmccc.rte.
```

Workaround: You may ignore these messages. The /etc/VRTSvcs/conf/config/ClusterConnectorConfigType.cf file is restored back to the original 5.0GA version.

This error is specific to the VRTScmccc.rte file set. The VRTScmcs.rte file set does not exhibit any errors when rejected using the installp -r command.

Warning messages in the log file when running the installmp command

The installmp command logs the following messages. [782798]

```
warning: user vcsbuild does not exist - using root
warning: group fcf does not exist - using root
warning: user vcsbuild does not exist - using root
warning: group fcf does not exist - using root
```

Workaround: None. You may ignore these messages.

Platform attribute in the ClusterConnector.config file is not updated

The Platform attribute in the ClusterConnector.config file remains set to Solaris irrespective of the installation platform. The ClusterConnector.config file is created by the ClusterConnectorConfig agent and is used to set values in resource type definitions and main.cf configurations for the agent. [837685]

The ClusterConnectorVersion attribute might have no value because this value is not used in the current release.

Known issue for the Migrate Site task

The Migrate Site task starts the Migrate Site wizard that enables you to migrate one or more global service groups to one or more remote target clusters. The Cluster Management Console does not migrate a faulted service group. If you attempt to migrate a faulted service group, you may see an entry similar to the following in the management server log:

```
2006-11-20 10:38:33 INFO    Unable to use the -force option when
the cluster that has Authority for the group is not completely
down {vrts.vxcs.mcm.gui.web.actions.wizard.MigrateSiteLastPage
lookupFinish()
```

Workaround: In the Global Application Group Selection panel, select only service groups that are in the online or partial state. Do not select service groups that are in the faulted state.

Erroneous output from gares command

The gares command returns a value for the Start attribute that is different from what the hares command returns. The local values are inverted (exchanged). For example, if gares returns 1, hares returns 0. [853969]

Workaround: This situation can result if the attribute values with local scope are missing for a newly-added system in the system list of a service group. Use the switch command for the CMC_CC service group (for configurations that use the cluster connector) or reconnect to the cluster (for configurations that use direct connection).

Cluster Management Console displays fatal errors

CMC displays fatal errors when it encounters an invalid XML file for an agent. [595973]

Workaround: None. Make sure the XML files for custom agents are valid.

The database fails to back up or restore to a path with Japanese characters

The database fails to back up or restore to the specified path with Japanese characters in it, when the command gadb -backup is run. [767796]

Workaround: Use English folder names when backing up, then copy the database file to the Japanese folder manually, if required.

Cannot uninstall updates on Windows management server

On Windows, uninstalling the VCS 5.0 MP1 management server using Add or Remove Programs removes only the entry from the Add or Remove Programs list. No files are removed. You must perform a management server uninstallation using the original VCS 5.0 uninstallation program. You cannot revert a VCS 5.0 MP1 management server back to a VCS 5.0 management server. [841149]

View displays incorrect version

After upgrading to the Cluster Management Console for VCS 5.0 MP1, the Admin:Management Server view (Admin –> Management Server) shows an incorrect version of 5.0.1136.0 and an incorrect installation history. The correct information is in the About box. [856103]

Default SMTP and SNMP addresses in notification policies for Cluster Management Console

When you configure notification settings, the Edit SMTP Settings task asks you to provide default email or default SNMP console addresses. The policy configuration wizard uses these addresses only to populate the recipient lists during policy configuration. The wizard does not automatically configure policies with these addresses.

When you launch the Notification Policy Configuration wizard, the default email address you specified appears in the Notification Recipients dialog box.

If you add email addresses to this list, the wizard adds them to the policy along with the default address. However, if you delete all the addresses from the Email Recipients list, including the default email address, the wizard configures no email addresses in the policy.

Leave default email addresses in the recipients list to configure them into the policy.

The same behavior applies to specifying default SNMP addresses.

Console displays logs in English and Japanese

If your management server is configured to run in the Japanese locale, but the managed cluster does not have the Japanese language pack installed, the management server displays a mix of logs in English and Japanese. [778176]

Workaround: Make sure the managed cluster has the Japanese language pack installed.

Some Cluster Management Console controls not immediately active

In some versions of Internet Explorer, you may need to click Flash-based screens, popups, and wizards once before the controls become active. Controls that require this activating click show the following message when you roll over them with your mouse pointer [603415]:

```
Press SpaceBar or Click to activate this Control
```

Login screen may not display after inactivity timeout

If your Cluster Management Console is inactive and the session times out, your next action in the console should return you to the login screen. However, if your next action is to request a sort or a new page, the console will not sort the data or load the page.

Workaround: Use the browser refresh feature and the login screen will display.

Very large clusters may not load into Cluster Management Console

Very large clusters may not load into Cluster Management Console. [493844]

Workaround: To accommodate very large clusters, increase the value of the loadClusterQueryTimeout property in the management server configuration file, /opt/VRTScmc/conf/ManagementServer.conf. The management server generates this file upon startup.

1. Stop the Cluster Management Server web console:

   ```
   /opt/VRTSweb/bin/stopApp cmc
   ```

2. Add the following line to the file /opt/VRTScmc/conf/ManagementServer.conf:

   ```
   loadClusterQueryTimeout=60000
   ```

   Adjust the value as needed to allow complete initial load of your cluster information.

3. Start the Cluster Management Server web console:

   ```
   /opt/VRTSweb/bin/startApp cmc ../VERITAS
   ```

Log entries in the Management Server:Logs view

The Management Server:Logs view might contain log entries for the management server and for the cluster. [610333]

Management server log entries have the value **`site`** in the Object Type column. Cluster log entries have the value **`cluster`** in the Object Type column.

Cannot install if VxAT 4.3 is installed

If you have installed Symantec Product Authentication Services on a system using the 4.3 client/server installer, install of Cluster Management Console will not succeed because the path to the AT binaries is not in the path. Since this path is not present, the custom action DLL in our MSI will not be able to run certain AT-related commands. [617861]

Workaround: Add the path for the AT binaries before attempting a Cluster Management Console install.

Uninstall of Cluster Connector in a secure cluster leaves the VxSS service group frozen

On UNIX, when you remove the cluster connector from a secure cluster, the VxSS service group is frozen. [619106]

Workaround: Manually unfreeze the VxSS group. Run the following commands.

```
haconf -makerw
hagrp -unfreeze VxSS -persistent
haconf -dump -makero
```

Windows management server uninstall using Add or Remove Programs does not remove folder

After using Add or Remove Programs to remove (uninstall) the Windows management server, an empty Cluster Management Console folder remains:

The default path is C:\Program Files\VERITAS.

Workaround: Delete the empty folder after the uninstall.

Windows cluster monitor uninstall does not remove folder

After a Windows cluster monitor uninstall, an empty folder remains:

The default path is C:\Program Files\VERITAS.

Workaround: Delete the empty folder after the uninstall.

Uninstalling Cluster Connector does not remove entry from Add\Remove Programs on Windows

After you uninstall cluster connector on Windows cluster nodes, the Add or Remove Programs control panel continues to show an entry for cluster connector. This persistent entry prevents any reinstallation of cluster connector. [599424]

Workaround: Remove the Veritas Cluster Management Console entry from the list using Windows Installer Cleanup Utility. Run the utility to remove the entry on each node. If you do not have the utility, you may download it from the Microsoft support site.

Windows install over Terminal Services needs Service Pack 4

Per Microsoft, Windows 2000 without at least Service Pack 4 has problems installing multiple MSI files that alter the same registry key over Terminal Services.

Workaround: If you want to install to a Windows 2000 host using Terminal Services, first ensure that the system has Windows 2000 Service Pack 4 installed.

Removing the *CMC_SERVICES* domain

Uninstalling the management server in multi-cluster environments does not remove the *CMC_SERVICES* domain. [612176]

You can verify the existence of this domain using the following command:

```
vssat showpd --pdrtype ab --domain CMC_SERVICES
```

You must manually remove the CMC_SERVICES domain using the command line. To manually remove all the peripherals in the CMC_SERVICES domain, enter the following command:

```
vssat deleteprpl --pdrtype ab --domain CMC_SERVICES --prplname
principalname
```

Enter the following command to remove the domain:

```
vssat deletepd --pdrtype ab --domain CMC_SERVICES@hostname
```

You can determine the host name using the following command:

```
vssat showpd
```

Other known issues

The -s option with the haping command does not work as expected

The `haping -s` command does not work as expected. [1211284]

VCS Simulator does not start on Windows systems

On Windows systems, starting VCS Simulator displays an error that the required MSVCR70.DLL is not found on the system. [859388]

Workaround: Run the following command:

```
set PATH=%PATH%;%VCS_SIMULATOR_HOME%\bin;
```

Or append %VCS_SIMULATOR_HOME%\bin; to PATH env variable.

Documentation errata

This section adds or replaces content in the VCS 5.0 documents.

Veritas Cluster Server User's Guide

User's Guide does not mention backward-compatibility of the Java Console

The VCS User's Guide does not mention the backward-compatibility of Cluster Manager (Java 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 a pre-4.1 version of a VCS cluster. [641680]

Updated definition of the IntentOffline attribute

The definition of IntentOnline needs to be updated to include following information:

VCS sets IntentOnline attribute value to 2 for failover groups while VCS attempts to autostart a service group. Once the service group is online, VCS sets IntentOnline value to 1. [831858]

Veritas Cluster Server Centralized Management Guide

This information replaces the information in the Veritas Cluster Server Centralized Management Guide for 5.0. Numbers in parentheses indicate the page number of the Centralized Management Guide where this information appears.

Backing up the database

Backing up the database (page 158) is necessary so that crucial configuration and historical information can be recovered in the event of a failure. You can back up the database using the Cluster Management Console or the CLI. During the backup task, an archived copy of the database file and the associated transaction log file are backed up to a physically separate location. This location can be a tape drive or a disk drive. [703139]

To backup the database to a file

1. In the **Administration: Management Server Database** view, in the **Operations** task panel, click **Backup database to a file**.
2. In the **Enter a valid directory or tape drive on the server** dialog box, enter an existing directory path on the management server.

   If the directory path you specify does not exist, the database backup command does not create it.
3. Click **OK**.

To backup the database to a file using the command line

◆ `gadb -backup -to archive`

This command creates an archive backup file that contains both the database and transaction log. The database archive file is created in the directory path specified by *archive*. The database archive file name is of the form:

```
CCAvailDbBackUp@yyyy-mm-dd_hh_mi_ss.1
```

The timestamp portion is in GMT.

Creating custom reports

The section on accessing the database information contains references to $ms_host, which is a variable. Read $ms_host as *ms_host*.

When configuring ODBC, replace *ms_host*. with the name of the management server host. Do not include the $ sign in the host name.

Veritas Cluster Server Bundled Agents Reference Guide

DNS agent

This information replaces the information in the guide.

The DNS agent updates and monitors the mapping for the following:

- the host name to IP address (A, AAAA, or PTR record)
- the canonical name (CNAME)

The agent performs these tasks for a DNS zone when failing over nodes across subnets (a wide-area failover). Resource records (RR) can include different types: A, AAAA, CNAME, name server, SOA, and PTR records.

Use the DNS agent when the failover source and target nodes are on different subnets. The agent updates the name server and allows clients to connect to the failed over instance of the application service.

Agent functions

Online — Sends a DNS query to retrieve the Start of Authority (SOA) record of the zone that the Domain agent attribute defines. The master server's name is in the SOA field. Unless you define the StealthMasters attribute, it is the only server for the update. When you define the StealthMasters attribute, only the servers that the attribute defines are updated.

The agent creates PTR records for each RR of type A or AAAA if the value of the CreatePTR attribute is true. A prerequisite for this feature is that the same master or stealth servers serve the forward (A or AAAA) and reverse zones.

Offline — If attribute OffDelRR is true, offline removes all records that the ResRecord keys define.

Monitor — Returns the ONLINE state if at least one name server reports all mappings that ResRecord defines. The name servers are the master or StealthMaster, and all the servers for which an NS record for the zone exists.

Clean — Removes the Online lock file, if it exists.

Open — Removes the Online lock file if the resource is reported online on another node inside the cluster to prevent concurrency violation. If the lock file exists, at least one name server has to report all the RRs that the ResRecord attribute defines. If one name server cannot report all the RRs, the agent function removes the Online lock file.

State definitions

ONLINE All the RRs that one name server reports.

OFFLINE Indicates an offline state when either of the following is true:

- The online lock does not exist.
- At least one server cannot report all of the RRs' mappings.

UNKNOWN A problem exists with the configuration.

Attributes

Table 1-3 Required attributes

Required attribute	Description
Domain	A string representing the DNS zone that the agent administers. The domain name can only contain alphanumeric symbols and the dash. Type and dimension: string-scalar Examples: ■ "demo.symantec.com" (forward mapping) ■ "2.168.192.in-addr.arpa" (IPv4 reverse mapping)

Table 1-3 Required attributes

Required attribute	Description
ResRecord	An association of DNS resource record values. Each ResRecord attribute consists of two values: *DNS record key = DNS record data*. Note that the record key must be a unique value. Type and dimension: association-scalar Examples: - For forward mapping, where the zone is demo.symantec.com: - sles901 = "192.168.2.191" - ww2 = sles901 - For forward mapping, where the zone is demo.symantec.com. A multi-home DNS record, typically for one host with two network interfaces, different address, but the same DNS name. This results in two-A records, or a single A record with continuation lines. sle902 = "192.168.2.102 10.87.13.22" - For reverse IPv4 address mapping, where the zone is 2.168.192.in-addr.arpa: 191 = "sles901.demo.symantec.com" - Use only partial host names. If you use a fully qualified domain name, append a period "." at the end of the name. For CNAME records, use: - ResRecord = { www = mydesktop } or - ResRecord = { www = "mydesktop.marketing.db.com." } Where the Domain attribute is "marketing.db.com" The agent uses case-insensitive pattern matching—and a combination of the Domain and ResRecord attribute values—to determine the resource record type. The RR type is as follows: - PTR: if the Domain attribute ends with .arpa - A: if the record data field is four sets of numbers, where a space separates each set. The following details the pattern it tries to match: [*1-223*].[*0-255*].[*0-255*].[*0-255*] Hexadecimal is not supported. - AAAA: if the record data fields are in multiple sets of hexadecimal format, then this record is an IPv6 associated type AAAA record. - CNAME: for any other results. **Note:** If a name in the ResRecord attribute does not comply with RFC 1035, then a warning is issued to the log file. The ResRecord association is not used.

Table 1-4 Optional attributes

Optional attribute	Description
TTL	A non-zero integer represents the “Time To Live” value, in seconds, for the DNS entries in the zone that you want to update. A lower value means more hits on your DNS server, while a higher value means more time for your clients to learn about changes. The time-in-seconds value may take the value 0, which indicates never caching the record, to a maximum of 2,147,483,647, which is over 68 years! The current best practice recommendation (RFC 1912) proposes a value greater than one day, and on RRs that do not change often, consider multi-week values. Type and dimension: integer-scalar Default: 86400 Example: "3600"
StealthMasters	The list of primary master name servers in the domain. This attribute is optional since the first name server is retrieved from the zones SOA (Start of Authority) record. If the primary master name server is a stealth server, define this attribute. A stealth server is a name server that is authoritative for a zone, but does not appear in zone’s SOA record. It is hidden to prevent direct attacks from the Internet. Type and dimension: string-keylist Example: { "10.190.112.23" }
TSIGKeyFile	Required when you configure DNS for secure updates. Specifies the absolute path to the file containing the private TSIG (Transaction Signature) key. Type and dimension: string-scalar Example: /var/tsig/Kexample.com.+157+00000.private

Table 1-4 Optional attributes

Optional attribute	Description
CreatePTR	Use the CreatePTR attribute to direct the online agent function to create PTR records for each RR of type A or AAAA. You must set the value of this attribute to true (1) to create the record. Before you can use this attribute, the same master or stealth servers must serve the forward (A or AAAA) and reverse zones. Type and dimension: boolean-scalar Default: 0 Example: 1
OffDelRR	Use the OffDelRR attribute to direct the offline agent function to remove all the records that the ResRecord key defines. You must set the value of this attribute to true (1) to have the agent remove all the records. The online agent function always adds records if they do not exist. Type and dimension: boolean-scalar Default: 0 Example: 1

Resource type definition

```
type DNS (
    static str ArgList[] = { Domain, TTL, TSIGKeyFile,
    StealthMasters, ResRecord, CreatePTR, OffDelRR }
    str Domain
    int TTL = 86400
    str StealthMasters[]
    str TSIGKeyFile
    str ResRecord{}
    boolean CreatePTR = 0
    boolean OffDelRR = 0
)
```

Monitor scenarios

Depending on the existence of the Online lock file and the defined Resource Records (RR), you get different status messages from the Monitor function.

Table 1-5 Monitor scenarios for the Online lock file

Online lock file exists	Expected RR mapping	Monitor returns
NO	N/A	OFFLINE
YES	NO	OFFLINE
YES	YES	ONLINE

Sample Web server configuration

Take the former Veritas corporate web server as an example. A person using a web browser specifies the URL www.veritas.com to view the Veritas Web page. Where www.veritas.com maps to the canonical name mtv.veritas.com, which is a host in Mountain View running the web server. The browser, in turn, retrieves the IP address for the web server by querying the domain name servers. If VCS fails the web server for www.veritas.com from Mountain View to Heathrow, the domain name servers must be updated with the new canonical name mapping. This update occurs so that the web browsers are directed to Heathrow instead of Mountain View. The DNS agent should update the name server to change the mapping of www.veritas.com. From mtv.veritas.com to the canonical name of the standby system in Heathrow, hro.veritas.com, in case of a failover.

Secure DNS update for BIND 9

The DNS agent expects that the zone's allow-update field contains the IP address for the hosts that can dynamically update the DNS records. This functionality is default for the DNS agent. Since a competent black hat can, however, spoof IP addresses, consider TSIG as an alternative.

TSIG (Transaction Signature) as specified in RFC 2845, is a shared key message authentication mechanism, which is available in DNS. A TSIG key provides the means to authenticate and verify the validity of exchanged DNS data. It uses a shared secret key between a resolver and either one or two servers to provide security.

Setting up secure updates using TSIG keys for BIND 9

In the following example, the domain is example.com.

To use secure updates using TSIG keys

1 Run the `dnssec-keygen` command with the HMAC-MD5 option to generate a pair of files that contain the TSIG key:

```
# dnssec-keygen -a HMAC-MD5 -n HOST example.com.
    Kexample.com.+157+00000
```

2 Open the Kexample.com.+157+00000.key file. After you run the `cat` command, the contents of the file resembles:

```
# cat Kexample.com.+157+00000.key
    example.com. IN KEY 512 3 157 +Cdjlkef9ZTSeixERZ433Q==
```

3 Copy the shared secret (the TSIG key), which looks like:

`+Cdjlkef9ZTSeixERZ433Q==`

4 Configure the DNS server to only allow TSIG updates using the generated key. Open the named.conf file and add these lines.

```
key example.com. {
    algorithm hmac-md5;
    secret "+Cdjlkef9ZTSeixERZ433Q==";
};
```

Where **`+Cdjlkef9ZTSeixERZ433Q==`** is the key.

5 In the named.conf file, edit the appropriate zone section and add the allow-updates sub-statement to reference the key:

```
allow-update { key example.com. ; } ;
```

6 Save and restart the named process.

7 Place the files containing the keys on each of the nodes that is listed in your group's SystemList. The DNS agent uses this key to update the name server.

Copy both the private and public key files on to the node. A good location is in the /var/tsig/ directory.

8 Set the TSIGKeyFile attribute for the DNS resource to specify the file containing the private key.

```
DNS www (
Domain = "example.com"
ResRecord = {www = north}
TSIGKeyFile a= "/var/tsig/Kexample.com.+157+00000.private"
)
```

NIC agent

This information is in addition to the NetworkHosts attribute in the guide.

NetworkHosts is a required attribute for virtual NIC devices.

Veritas Cluster Server Installation Guide

Supplemental: Using a Switch or Interface in a Virtual I/O Environment Requires Configuration Changes

In a virtual I/O environment, the kernel and the interface card or switch use different maximum transfer unit (MTU) values. If a mismatch exists, packet loss for larger packets can result.

LLT receives its MTU value from the kernel. LLT queries the AIX native DLPI layer to get the MTU size for each private network interface. It then takes the least of all the interface MTU values. LLT uses this least value as an overall MTU value, and uses this value for communication with its peers.

See Veritas technote 278286 for more information.

Example:

If one private link has jumbo frames and the other has normal MTU, then the overall MTU is normal (1,500 bytes). Use the command lltstat -c to get the overall MTU value:

```
# lltstat -c | grep overall
overall mtu: 1460
```

If the LLT private network links are virtual Ethernet devices, the AIX DLPI layer returns an MTU size of 65,354 bytes for each link. The overall MTU, as a result, is 65,354 bytes. This value does not match with the maximum MTU that the external switch can handle. As a result, the switch drops any LLT packets greater than 1,500 bytes.

Workaround 1:

If the LLT private links are over virtual Ethernet devices, modify the /etc/llttab entry to restrict the MTU value to 1500.

Sample llttab file restricting the MTU size to 1500:

```
# more /etc/llttab
set-node nodeamp8
set-cluster 90
link en1 /dev/dlpi/en:1 - ether - 1500
link en2 /dev/dlpi/en:2 - ether - 1500
```

Workaround 2:

Connect the interfaces with crossover cables instead of a switch. This is restrictive however. [430935]

Symantec Product Authentication Service 4.3.x required for cluster connector installation

You must install cluster connector from a system that has Symantec Product Authentication Service 4.3.x, or at least the authentication broker installed.

You can also install cluster connector from a cluster node, provided that you are installing cluster connector on nodes that are part of the same cluster. [611353]

Cluster connector must be taken offline before it is uninstalled

Ensure that you take the CMC service group offline before you uninstall cluster connector. Otherwise, cluster connector remains running even after you uninstall the cluster connector software. [796739]

To take the CMC service group offline on UNIX platforms

1. Obtain a command prompt on the management server host system.
2. Enter the following command:

   ```
   hagrp -offline CMC -sys
   ```

   Replace sys with the name of the system that is running the CMC service group.

Veritas High Availability Agent for DB2 Installation and Configuration Guide

In the Veritas High Availability Agent for DB2 Installation and Configuration Guide, replace all custom file names of custom_monitor_$db2instance _$nodenum, with monitor_custom_$db2instance_$nodenum. [786209]

The section on Disabling in-depth monitoring refers to MonScript when it should refer to IndepthMonitor (786221).

Veritas High Availability Agent for Oracle Installation and Configuration Guide

On page 14, the command documented to invoke the Info entry point is not correct. Use the following command:

```
hares -value resource ResourceInfo [system]\
[-clus cluster | -localclus]
```

On page 18, the description of IGNORE action is missing the following information:

When the Veritas Agent for Oracle encounters an error that does not have a matching error code in the oraerror.dat file, then the agent ignores the error.

Software limitations

The following limitations apply to this release.

HAD reports incorrect CPU utilization

API issues cause HAD to report incorrect CPU use. [1214140]

Cluster address for global cluster requires resolved virtual IP

The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.

System names in VCS

Systems specified in the VCS configuration file, main.cf, and in the files /etc/nodename and /etc/llthosts, must be consistent. The names cannot include periods and thus must not be in the fully qualified form. If you create the file /etc/VRTSvcs/conf/sysname to contain system names used by main.cf, VCS uses the file to verify the names.

Systems in a cluster must have same system locale setting

VCS does not support clustering of systems with different system locales. All systems in a cluster must be set to the same locale.

GAB panics the systems while VCS gets diagnostic data

On receiving a SIGABRT signal from GAB, VCS engine forks off `vcs_diag` script. When VCS engine fails to heartbeat with GAB, often due to heavy load on the system, the `vcs_diag` script does a `sys req` to dump the stack trace of all processes in the system to collect diagnostic information. The dump of stack trace is intended to give useful information for finding out which processes puts heavy load. However, the dumping puts extra load on the system that causes GAB to panic the system in such heavy loads. See *VERITAS Cluster Server User's Guide* for more information.

Workaround: Disable the `vcs_diag` script. To disable, rename the file /opt/VRTSvcs/bin/vcs_diag to /opt/VRTSvcs/bin/vcs_diag.backup.

Using agents in NIS

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can hang if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to hang and possibly time out. For example, if the host is configured to use NIS maps as a client, basic commands such as `ps -ef` can hang if there is network disconnect. Symantec recommends creating users locally and configuring /etc/netsvc.conf to reflect local users.

Fire drill does not support volume sets

The fire drill feature for testing fault readiness of a VCS configuration supports only regular Volume Manager volumes. Volume sets are not supported in this release.

Manually removing VRTSat package erases user credentials

Symantec recommends saving user credentials before manually removing the VRTSat package. If you need the credentials again, you can restore them to their original locations.

To save user credentials

1 Run the `vssat showbackuplist` command. The command displays the data files and backs them up into the SnapShot directory /var/VRTSatSnapShot. Output resembles the following:

```
 vssat showbackuplist
B| /var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B| /var/VRTSat/.VRTSat/profile/certstore
B| /var/VRTSat/RBAuthSource
B| /var/VRTSat/ABAuthSource
B| /etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapShot
```

2 Move the credentials to a safe location. Preserving the directory structure makes restoring the files easier.

To restore user credentials

1 Navigate to the SnapShot directory or the safe location where you previously saved credentials:

```
cd /var/VRTSatSnapShot/profile
```

2 Restore the files:

```
cp ABAuthSource /var/VRTSat/
cp RBAuthSource /var/VRTSat
cp VRTSat.conf /etc/vx/vss
cd /var/VRTSatSnapShot/
cp -r profile /var/VRTSat/.VRTSat
```

Limitation with RDAC driver and FAStT array for coordinator disks that use raw disks

For multipathing to connected storage, AIX uses the RDAC driver for FAStT arrays. Since it is an active/passive array, only the current active path is exposed to clients. The I/O fencing driver, vxfen, can use only a single active path and has no foreknowledge of the passive paths to the coordinator disks on an array. If the single active path fails, all nodes in the cluster lose access to the coordinator disks.

The loss of the path to the coordinator disks can potentially go unnoticed until a reboot, split brain, or any other reason that leads to a cluster membership change occurs. In any of these conditions, the cluster cannot form, and all nodes panic to prevent data corruption. No data loss occurs.

Workaround: Use DMP and specify paths to coordinator disks as DMP paths rather than raw disks to avoid this limitation.

I/O fencing limitations

Fencing is not supported in a VIO server environment

Certain SCSI3-PR command subsets that are critical to use fencing are not yet available.

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *Veritas Cluster Server User’s Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

I/O fencing uses SCSI-III Persistent Reserve keys to implement data protection. Keys are placed on I/O fencing coordinator disks and on data disks. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator disks and data disks to prevent possible difficulties with subsequent cluster startup.

Using the `reboot` command rather than the `shutdown` command bypasses shutdown scripts and can leave keys on the coordinator disks and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the `shutdown -r` command on one node at a time and wait for each node to complete shutdown.

Virtualizing shared storage using VIO servers and client partitions

AIX 5.3, with proper patches to the operating system and client partitions, is capable of running multiple virtualized partitions within a single frame. You can split the CPU, memory, and certain adapters (networking and storage), into smaller virtual units that the partitions can then use.

In an Advanced POWER™ Virtualization (APV) environment, AIX uses the VIO Server to monitor and manage the I/O paths for the virtualized client partitions. At a very high level, the VIO server provides a partition’s access to storage that is external to the physical computer. The VIO server encapsulates the physical hardware into virtual adapters called virtual SCSI adapters (server adapter). On the client side, you can create virtual adapters (client adapters) that map to the server adapter and enable a partition to connect to external storage.

Note: Fencing and the LVMVG agent are not supported in a VIO server environment.

The VIO server provides similar mechanisms to share limited networking resources across partitions. Refer to the manual that came with your system to help set up partitions, and to configure and use the various components such as VIO server and HMC, which are integral parts of IBM's APV environment.

The minimum patch level for using VIO servers with VCS is: Fix Pack 7.1.2.0.0.

Supported Storage

Refer to the IBM data sheet:
http://techsupport.services.ibm.com/server/vios/documentation/datasheet.html

Disk Restrictions

When using VCS in combination with VIO servers and their client partitions, you need to ensure that no reservations are placed on the shared storage. This enables client partitions on different systems to access and use the same shared storage.

- If the shared storage is under MPIO control, set the reserve_policy attribute of the disk to no_reserve.
- If the shared storage is not under MPIO control, look up the array documentation to locate a similar attribute to set on the disk.

Internal testing on EMC disks shows that this field maps as the reserve_lock attribute for EMC disks. In this case, setting it to `no` achieves the same result.

Accessing the same LUNs from Client Partitions on Different Central Electronics Complex (CEC) Modules

This section briefly outlines how to set shared storage so that it is visible from client partitions on different CEC modules.

With the VIO server and client partitions set up and ready, make sure that you have installed the right level of operating system on the client partitions, and that you have mapped the physical adapters to the client partitions to provide access to the external shared storage.

To create a shareable diskgroup, you need to ensure that the different partitions use the same set of disks. A good way to make sure that the disks (that are seen from multiple partitions) are the same is to use the disks serial numbers, which are unique.

Run the following commands on the VIO server (in non-root mode), unless otherwise noted.

Get the serial number of the disk of interest:

```
$ lsdev -dev hdisk20 -vpd
hdisk20
U787A.001.DNZ06TT-P1-C6-T1-W500507630308037C-
L4010401A00000000  IBM FC 2107

Manufacturer................IBM
Machine Type and Model......2107900
Serial Number...............7548111101A
EC Level....................131
Device Specific.(Z0)........10
Device Specific.(Z1)........0100
…
```

Make sure the other VIO server returns the same serial number. This ensures that you are viewing the same actual physical disk.

List the virtual SCSI adapters.

```
$ lsdev -virtual | grep vhost
vhost0  Available  Virtual SCSI Server Adapter
vhost1  Available  Virtual SCSI Server Adapter
```

Note: Usually vhost0 is the adapter for the internal disks. vhost1 in the example above maps the SCSI adapter to the external shared storage.

Prior to mapping hdisk20 (in the example) to a SCSI adapter, change the reservation policy on the disk.

```
$ chdev -dev hdisk20 -attr reserve_policy=no_reserve
    hdisk20 changed
```

For hdisk20 (in the example) to be available to client partitions, map it to a suitable virtual SCSI adapter.

If you now print the reserve policy on `hdisk20` the output resembles:

```
$ lsdev -dev hdisk20 attr reserve_policy
value
no_reserve
```

Next create a virtual device to map `hdisk20` to `vhost1`.

```
$ mkvdev -vdev hdisk20 -vadapter vhost1 -dev mp1_hdisk5
mp1_hdisk5 Available
```

Finally on the client partition run the `cfgmgr` command to make this disk visible via the client SCSI adapter.

You can use this disk (hdisk20 physical, and known as mp1_hdisk5 on the client partitions) to create a diskgroup, a shared volume, and a eventually a shared file system.

Perform regular VCS operations on the clients vis-a-vis service groups, resources, resource attributes, etc.

Using a switch or interface in a virtual I/O environment requires configuration changes

In a virtual I/O environment, the kernel and the interface card or switch use different maximum transfer unit (MTU) values. If a mismatch exists, packet loss for larger packets can result.

LLT receives its MTU value from the kernel. LLT queries the AIX native DLPI layer to get the MTU size for each private network interface. It then takes the least of all the interface MTU values. LLT uses this least value as an overall MTU value, and uses this value for communication with its peers.

See Veritas technote 278286 for more information.

Example:

If one private link has jumbo frames and the other has normal MTU, then the overall MTU is normal (1,500 bytes). Use the command `lltstat -c` to get the overall MTU value:

```
 lltstat -c  | grep overall
overall mtu: 1460
```

If the LLT private network links are virtual Ethernet devices, the AIX DLPI layer returns an MTU size of 65,354 bytes for each link. The overall MTU, as a result, is 65,354 bytes. This value does not match with the maximum MTU that the external switch can handle. As a result, the switch drops any LLT packets greater than 1,500 bytes.

Workaround:

- If the LLT private links are over virtual Ethernet devices, modify the /etc/llttab entry to restrict the MTU value to 1500.

  Sample llttab file restricting the MTU size to 1500:

  ```
   more /etc/llttab
  set-node nodeamp8
  set-cluster 90
  link en1 /dev/dlpi/en:1 - ether - 1500
  link en2 /dev/dlpi/en:2 - ether - 1500
  ```

  Or

- Connect the interfaces with crossover cables instead of a switch. This is restrictive however.

Bundled agent limitations

Volume agent clean may forcibly stop volume resources

When the attribute FaultOnMonitorTimeouts calls the Volume agent `clean` entry point after a monitor time-out, the `vxvol -f stop` command is also issued. This command forcibly stops all volumes, even if they are still mounted.

NFS failover

If the NFS share is exported to the world (*) and the NFS server fails over, NFS client displays the following error, “Permission denied”.

To avoid this error, export NFS shares explicitly using FQDN hostnames.

False concurrency violation when using PidFiles to monitor application resources

The PID files created by an application contain the PIDs for the processes that are monitored by Application agent. These files continue to exist even after a node running the application crashes. On restarting the node, the operating system may assign the PIDs listed in the PID files to other processes running on the node.

Thus, if the Application agent monitors the resource using the PidFiles attribute *only*, the agent may discover the processes running and report a false concurrency violation. This could result in some processes being killed that are not under VCS control.

Networking agents do not support IPv6 protocol

The bundled IP, NIC, IPMultiNIC, MultiNICA, IPMultiNICB, and MultiNICB agents for VCS 5.0 do not support the IPv6 enhanced IP protocol.

VCS does not provide a bundled agent for volume sets

VCS 5.0 does not provide a bundled agent to detect Volume Manager volume sets, Problems with volumes and volume sets can only be detected at the DiskGroup and Mount resource levels.
Workaround: Set StartVolumes and StopVolumes attributes of the DiskGroup resource that contains volume set to 1. If a file system is created on the volume set, use a Mount resource to mount the volume set.

No LVMVG agent support in a VIO environment

The LVMVG agent is not supported in a VIO environment.

Cluster Management Console limitations

Cluster connector not supported on some OS versions

Cluster Management Console does not support cluster connector on AIX 5.1, Solaris 7, and RHEL 3.0. If your cluster runs on any of these platforms, you must use direct connection to manage the cluster from a management server.

Limited peer management server support

Peer management server support is limited to a configuration of two management servers in an enterprise. An enterprise of three or more management servers is not supported in this release.

Management server cannot coexist with GCM 3.5 Master

The Cluster Management Console management server should not be installed on the same system with a GCM 3.5 Master. These two products will conflict with each other and are not supported running on the same system.

Agent info files needed for Agent Inventory report

By design, the Agent Inventory report requires agent info files that supply the information reported on individual agents. These files are shipped with agents in VCS.

Global clusters must be CMC-managed clusters

All clusters forming a global cluster (using the VCS 4.0 Global Cluster Option) must be managed clusters in order for Veritas Cluster Management Console views to display correct and consistent information. Managed clusters are running the cluster connector or have a direct connection with the management server.

Windows Active Directory installation requires NetBIOS

If you install Cluster Management Console management server in a Windows Active Directory domain, NetBIOS must be turned on. A native (non-NetBIOS) Active Directory environment is not supported in this release.

Remote root broker not supported on Windows

If you set up a management server on a Windows system, you must configure a root broker on the management server system. This release does not support specifying a remote root broker during management server install [841739].

The root broker can be changed after install using the configureRemoteRoot.exe installed in C:\Program Files\VERITAS\Cluster Management Console\bin (default install directory).

Cluster Manager (Java console) limitations

Use the VCS 5.0 Java Console to manage clusters

Cluster Manager (Java Console) from previous VCS versions cannot be used to manage VCS 5.0 clusters. Symantec recommends using the latest version of Cluster Manager. See the *Veritas Cluster Server 5.0 Installation Guide* for instructions on upgrading Cluster Manager.

Run Java Console on a non-cluster system

Symantec recommends not running Cluster Manager (Java Console) for an extended period on a system in the cluster. The Solaris version of the Java Virtual Machine has a memory leak that can gradually consume the host system’s swap space. This leak does not occur on Windows systems.

Cluster Manager and wizards do not work if the hosts file contains IPv6 entries

VCS Cluster Manager and Wizards fail to connect to the VCS engine if the /etc/hosts file contains IPv6 entries.

Workaround: Remove IPv6 entries from the /etc/hosts file.

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the UseFence attribute is set to the default, “None.”

Undocumented commands, command options, and libraries

VCS contains undocumented commands and command options intended for development use only. Undocumented commands are not supported.

Documentation

Product guides are available on the documentation disc in PDF and HTML formats. We recommend copying pertinent information, such as installation guides and release notes, from the disc to your system directory /opt/VRTS/docs for reference.

VCS documentation set

VCS includes the following documents.

Title	File Name
Veritas Cluster Server Installation Guide	vcs_install.pdf
Veritas Cluster Server Release Notes	vcs_notes.pdf
Veritas Cluster Server User's Guide	vcs_users.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev.pdf
Veritas Cluster Server Centralized Management Guide	vcs_central_mg.pdf
Veritas High Availability Agent for DB2 Installation and Configuration Guide	vcs_db2_install.pdf
Veritas High Availability Agent for Oracle Installation and Configuration Guide	vcs_oracle_install.pdf
Veritas High Availability Agent for Sybase Installation and Configuration Guide	vcs_sybase_install.pdf

The manual pages for the `VRTSllt`, `VRTSgab`, and `VRTSvcs` are installed in /opt/VRTS/man. Set the `MANPATH` environment variable so the `man`(1) command can point to the VCS manual pages.

For Bourne or Korn shell (sh or ksh), type:

```
MANPATH=$MANPATH:/opt/VRTS/man
export MANPATH
```

For C shell (csh or tcsh), type:

```
setenv MANPATH ${MANPATH}:/opt/VRTS/man
```

For more information, refer to the `man`(1) manual page.

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to clustering_docs@symantec.com.

Include the title and part number of the document (located in the lower left corner of the title page), and chapter and section titles of the text on which you are reporting.

Getting help

For technical assistance, visit
http://www.symantec.com/business/support/assistance_care.jsp

and select phone or email support. Select a product to use the Knowledge Base Search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and the customer email notification service. If you encounter an error when using a product, include the error number preceding the message when contacting Technical Services. You can also use the error number to search for information in TechNotes or documents on the website.