Symantec™ ApplicationHA User's Guide

Linux on KVM

6.0

December 2011

Symantec

Symantec™ ApplicationHA User's Guide

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

Product version: 6.0

Document version: 6.0.0

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our Web site at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Contents

Chapter 1

Introducing Symantec ApplicationHA

This chapter includes the following topics:

What is Symantec ApplicationHA

Symantec ApplicationHA provides monitoring capabilities for applications running inside guest virtual machines in the KVM virtualization environment. Symantec ApplicationHA adds a layer of application awareness to the core high availability (HA) functionality offered by Veritas™ Cluster Server (VCS) in the physical host.

Symantec ApplicationHA is based on VCS and uses similar concepts such as agents, resources, and service groups. However, it does not include the high availability cluster components such as the Group Membership and Atomic Broadcast (GAB), Low Latency Transport (LLT), Asynchronous Monitoring Framework (AMF), and Veritas Fencing (VxFEN). Symantec ApplicationHA has a lightweight server footprint that allows faster installation and configuration.

Key benefits include the following:

- Out of the box integration with VCS.

- Full visibility and control over applications with the ability to start, stop, and monitor applications running inside virtual machines.
- High availability of the application as well as the virtual machine inside which the application runs.
- Graded application fault-management responses such as:-
  - Application restart
  - ApplicationHA-initiated, graceful internal reboot (soft reboot) of a virtual machine
  - VCS-initiated, external reboot (hard reboot) of virtual machine
  - Failover of the virtual machine to another VCS node.
- Standardized way to manage applications using a single interface that is integrated with the Veritas Operations Manager (VOM) console.
- Specialized Application Maintenance mode, in which ApplicationHA allows you to intentionally take an application out of its purview for maintenance or troubleshooting.

How ApplicationHA is deployed in the KVM environment

Kernel-based Virtual Machine (KVM) is a full virtualization solution for Linux on AMD64 and Intel64 hardware. KVM lets you create and manage multiple virtual machines on a single physical host.

In the KVM virtualization environment, ApplicationHA provides high availability of applications running on virtual machines. Veritas Cluster Server (VCS) provides high availability of the virtual machines that run on the physical host.

The following figure illustrates how ApplicationHA and VCS are deployed in a typical KVM virtualization environment.

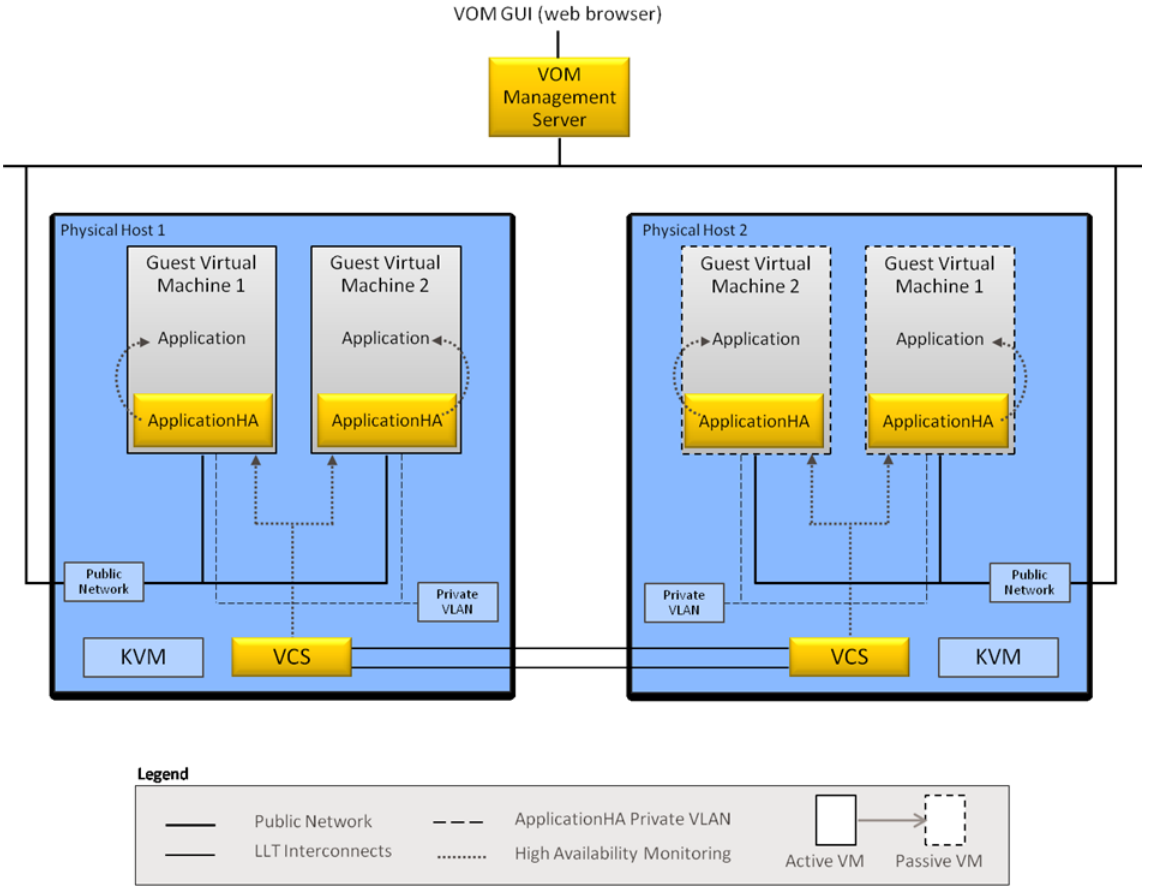

ApplicationHA is installed on the virtual machine, and provides high availability to a configured application running on the virtual machine. VCS is installed on the physical host, as part of a Storage Foundation Cluster File Server High Availability (SFCFSHA) stack installation. VCS provides high availability to the virtual machine where the configured application runs.

To ensure application-aware monitoring of virtual machines, you must enable VCS support for ApplicationHA.

See "Enabling VCS support for ApplicationHA" on page 30.

When you enable VCS to support ApplicationHA, a private VLAN is created between monitored virtual machines and the VCS node (physical host). The private VLAN facilitates heartbeat communication between VCS in the physical host and ApplicationHA in the virtual machines.

Veritas Operations Manager (VOM) provides you with a centralized management console (GUI) to administer application monitoring with ApplicationHA.

For more information on how VCS monitors virtual machines for high availability, see the *SFHA Virtualization Solutions Guide for Linux*.

How Symantec ApplicationHA works with VCS

Symantec ApplicationHA installed in virtual machines communicates directly with VCS installed in the physical host. Symantec ApplicationHA conveys the application health status to VCS in the form of a heartbeat. If VCS does not receive the heartbeat from a particular virtual machine within a specified interval, VCS either restarts that virtual machine or fails it over to another physical host.

You can monitor an application running on a virtual machine by using Veritas Operations Manager (VOM). An ApplicationHA add-on for VOM integrates the ApplicationHA tab with VOM.

The ApplicationHA tab is the primary interface for performing the application monitoring operations on a virtual machine. From this tab, you configure application monitoring and then monitor and control the configured application on the virtual machine. After configuring application monitoring, the ApplicationHA tab displays the state of the application and the component dependencies.

The following figure displays the ApplicationHA tab where Oracle is configured for monitoring.

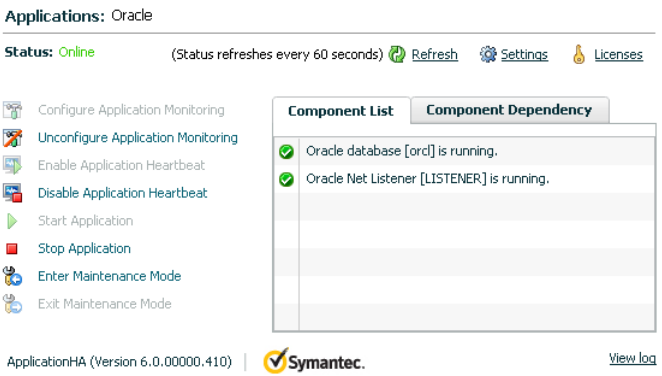

How Symantec ApplicationHA detects application failures

Symantec ApplicationHA architecture uses the agent framework to monitor the state of the applications and their dependent components running inside the

virtual machines. Symantec ApplicationHA agents monitor the overall health of the configured applications by running specific commands, tests, or scripts. For more details, see the agent functions section of the application-specific agent guides or the generic agent guide distributed with ApplicationHA.

The ApplicationHA Heartbeat agent is auto-configured in the virtual machine when you configure application monitoring. The Heartbeat agent sends the application heartbeat to VCS in the physical host. Symantec ApplicationHA uses the application heartbeat as the communication medium to convey the status of the application to VCS.

If an application fails, the application agents attempt to restart the application for a configurable number of times. If the agents are unable to start the application, ApplicationHA tries to reboot the virtual machine. After the virtual machine is restarted, Symantec ApplicationHA attempts to start the application and its dependent components in a predefined order.

Which applications can I monitor?

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

- Defined start, stop, and monitor procedures exist.
  The application to be monitored must have defined procedures for starting, stopping, and monitoring, as follows:

Start procedure	The application must have a command to start it and all the dependent components and resources it may require. Symantec ApplicationHA brings up the required resources in a specific order and then brings up the application using the defined start procedure.
Stop procedure	The application must have a command to stop it and all the dependent components and resources. Symantec ApplicationHA stops the application using the defined stop procedure, and then stops the required resources in an order that is reverse of the order in which they were started.

Monitor procedure
: The application must have a monitor procedure that determines if the specified application instance is healthy. The application must allow individual monitoring of unique instances. For example, in a database environment, the monitoring application can connect to the database server and perform SQL commands to verify read and write access to the database.

  The closer a test comes to matching what a user does, the better the test is in discovering problems. You should balance the level of monitoring between ensuring that the application is up and minimizing monitor overhead.

- Ability to restart the application in a known state
  When the application is stopped, it must close out all tasks, store data properly, and then exit. When Symantec ApplicationHA attempts to restart the application, it should be able to start from the last known state. In case of a server crash, the application must be able to recover gracefully.
  Commercial databases such as Sybase and Oracle are good examples of well-written, crash-tolerant applications. On any client request, the client is responsible for holding the request until it receives acknowledgement from the server. When the server receives a request, it is placed in a special redo log file. The database confirms that the data is saved before it sends an acknowledgement to the client.
  After a server crashes, the database recovers to the last-known committed state by mounting the data tables and applying the redo logs. This returns the database to the time of the crash. The client resubmits any outstanding client requests that are unacknowledged by the server, and all others are contained in the redo logs.

Components of the Symantec ApplicationHA setup

A Symantec ApplicationHA setup in the KVM virtualization environment comprises of the following components:.

- Symantec ApplicationHA guest components for virtual machines
- VCS in the virtualization infrastructure
- VOM add-on for Symantec ApplicationHA Management

Symantec ApplicationHA guest components for virtual machines

The Symantec ApplicationHA guest components are installed separately on the virtual machines where you wish to monitor applications. The guest components

include the configuration wizard and the ApplicationHA agents that are used for configuring and monitoring applications.

The guest components also include the Veritas Storage Foundation Messaging Service (xprtld). This service communicates the status of the applications running on the virtual machine and displays it in the ApplicationHA tab of the Veritas Operations Manager console.

VCS in the virtualization infrastructure

Veritas Cluster Server (VCS) is installed as part of a Storage Foundation Cluster File System High Availability (SFCFSHA) stack installation on a physical host. VCS is installed on more than one physical host to form a VCS cluster. As a result, VCS provides high availability in the infrastructure layer of the KVM virtualization environment on such physical hosts. VCS mainly ensures high availability of the virtual machines on which ApplicationHA monitors configured applications.

For more information on how ApplicationHA and VCS are integrated in the KVM virtualization environment:

See " How ApplicationHA is deployed in the KVM environment" on page 12.

For more information on how to install VCS as part of an SFCFSHA installation, see the *SFCFSHA Installation Guide*.

VOM add-on for Symantec ApplicationHA Management

The Veritas Operations Manager (VOM) Add-on for ApplicationHA Management is installed on the VOM Management Server. You must also add as managed hosts to VOM, the virtual machines where you want ApplicationHA to monitor applications. The ApplicationHA tab then appears on the VOM console for the respective virtual machine, and lets you administer application monitoring with ApplicationHA in the KVM environment.

Symantec ApplicationHA user privileges

Symantec ApplicationHA provides a set of privileges that are available when using VOM Console to manage ApplicationHA. These privileges define the application monitoring operations that a user can perform on the virtual machines. You can create roles and then assign privileges to the roles or assign privileges to the existing roles that are available in the virtualization environment. Application monitoring operations are enabled or disabled depending on the privileges that are assigned to the VOM user account. For example, the Admin privilege is required for configuring application monitoring on a virtual machine.

VOM administrators can use these privileges to configure access control in an application monitoring environment.

Symantec ApplicationHA provides the following privileges:

- View Application Monitoring State (Guest)
  Can view the application monitoring status on the virtual machine. The Guest cannot perform any ApplicationHA operations.
- Control Application Monitoring (Operator)
  Can perform all the ApplicationHA operations that include start and stop configured applications, enable and disable application monitoring, specify the application monitoring configuration settings, enter and exit application monitoring maintenance mode, and view application monitoring status.
  The Operator cannot configure or unconfigure application monitoring on the virtual machine.
- Configure Application Monitoring (Admin)
  Can perform all ApplicationHA operations that include configure and unconfigure application monitoring, start and stop configured applications, enable and disable application monitoring, specify the application monitoring configuration settings, enter and exit application monitoring maintenance mode, and view application monitoring status.

Symantec ApplicationHA agents

Agents are application-specific modules that plug into the ApplicationHA framework that manages applications and resources of predefined resource types on a system. The agents are installed when you install Symantec ApplicationHA guest components. These agents start, stop, and monitor the resources configured for the applications and report state changes. If an application or its components fail, ApplicationHA restarts the application and its resources on the virtual machine.

Symantec ApplicationHA agents are classified as follows:

- Infrastructure agents
  Agents such as NIC, IP, and Mount are classified as infrastructure agents. Infrastructure agents are automatically installed as part of the ApplicationHA installation on virtual machines.
  For more details about the infrastructure agents, refer to the *Veritas Cluster Server 6.0 Bundled Agents Reference Guide (Linux).*
- Application agents
  Application agents are used to monitor third party applications such as Oracle. These agents are packaged separately and are available in the form of an agent

pack that gets installed when you install Symantec ApplicationHA guest components.

The ApplicationHA agent pack is released on a quarterly basis. The agent pack includes support for new applications as well as fixes and enhancements to existing agents. You can install the agent pack on an existing ApplicationHA guest components installation.

Refer to the Symantec Operations Readiness Tools (SORT) Web site for information on the latest agent pack availability.

https://sort.symantec.com/agents

Refer to the agent-specific configuration guide for more details about the application agents.

About Symantec ApplicationHA licensing

Symantec ApplicationHA is a licensed product. Licensing for Symantec ApplicationHA is applicable for ApplicationHA guest components and is based on the server operating systems in use.

An evaluation license key is embedded in the product. This license key is valid only for a period of 2 months. If you are installing ApplicationHA for the first time, you can use the embedded license key or procure a permanent license key and enter the same while installing the product.

You can add or view the license keys from a virtual machine that has ApplicationHA guest components installed. You can add a license key through the command line or the ApplicationHA tab. For more information:

See "About managing ApplicationHA licenses" on page 57.

Chapter 2

Working with VOM

This chapter includes the following topics:

About Veritas Operations Manager

Veritas Operations Manager (VOM) provides you with a single, centralized management console for the Veritas Storage Foundation and High Availability products. You can use it to monitor, visualize, and manage storage resources and generate reports about them. VOM lets administrators centrally manage diverse datacenter environments.

A typical VOM deployment consists of a Management Server and the managed hosts. The managed host can be a physical or virtual system, that runs on any platform that VOM supports.

For more information on installing the VOM Management Server, see the *Veritas Operations Manager Installation Guide*.

Adding virtual machines and physical hosts to VOM

When you click a particular virtual machine on the Veritas Operations Manager (VOM) console, for the ApplicationHA tab to be visible, you must add the virtual machine as a managed host to VOM.

Symantec recommends that you also add as a managed hosts to VOM. the physical hosts that run the required virtual machine. This step helps you determine which

virtual machine is running on which physical host, especially when a failover occurs.

To add a virtual machine to VOM

1. On the Veritas Operations Manager console, click **Settings** > **Host Management**.
2. Click **Actions** > **Add Host(s)**.
3. On the **Add (Host)s** page, in the appropriate fields, specify the IP address, user name, and the password that you want to set for the managed host.
4. Click **Next**.

Note: You can use the same steps to add a physical host to VOM.

Accessing the ApplicationHA tab

To administer an application on a virtual machine that is running in the KVM environment, you must access the ApplicationHA tab of the Veritas Operations Manager (VOM) console.

In the ApplicationHA tab, you can perform administrative actions such as:

- Start an application
- Stop an application
- Configure application monitoring
- Unconfigure application monitoring
- Enable application heartbeat
- Disable application heartbeat
- Enter maintenance mode
- Exit maintenance mode

To access the ApplicationHA tab

1. On the Veritas Operations Manager console, click **Manage** > **Servers** > **Hosts**.
2. In the left pane, in the **License** list box, select the **ApplicationHA** check box.
3. In the right pane, click the virtual machine where you want to perform administrative actions.
4. Click the **ApplicationHA** tab.

Configuring Symantec ApplicationHA access control using VOM

The security or access control model in Veritas Operations Manager (VOM) is based on security groups. Security groups define privileges for user groups inside business entities, by using pre-defined user roles in VOM.

For example, Unices users can form a Unix security group, where the authentication mechanism is UNIX password authentication.

A business entity is a grouping construct for objects in VOM. The objects can be physical systems or virtual systems added to VOM as managed hosts. Using business entities, these objects are grouped for purposes such as reporting, alerting, and access control.

VOM defines three roles for access to ApplicationHA systems:

- Admin
- Operator
- Guest

For more details on the tasks and privileges of each role:

See "Symantec ApplicationHA user privileges" on page 17.

The following section broadly describes the steps to assign a role to a user. For the detailed steps, and for more information on security groups and business entities, see the *Veritas Operations Manager Administrators's Guide*.

To assign a role to a user

1 Create a business entity with one or more ApplicationHA systems.

2 Create a security group for the VOM user.

a. Specify the name of the security group. The name of the security group must exactly match the name of the group on the VOM Management Server, to which the user belongs.

b. From the available list of roles, select a role for the security group.

c. Select **Selected Business Entities** as scope of the role, and specify the business entity that you created for ApplicationHA systems in step 1.

Now any user that is part of the user group for which you configured a security group in step 2, can access the ApplicationHA systems with the defined role.

Chapter 3

Configuring application monitoring with Symantec ApplicationHA

This chapter includes the following topics:

About configuring application monitoring with Symantec ApplicationHA

ApplicationHA enables you to configure application monitoring for third party applications, in a virtualization environment.

For details refer to the respective agent configuration guide. You can download the guides from here:

https://sort.symantec.com/documents/

Consider the following before you proceed:

- You can configure application monitoring on a virtual machine using the Symantec ApplicationHA Configuration Wizard. The wizard is launched when you click **Configure Application Monitoring** on the ApplicationHA tab in Veritas Operations Manager (VOM).
- Apart from the application monitoring configuration, the configuration wizard also sets up the other components required for Symantec ApplicationHA to successfully monitor the applications.

You must first configure application monitoring using the configuration wizard before using VOM or VCS commands to add additional components or modify the existing configuration.

- You can use the wizard to configure monitoring for only one application per virtual machine.
  To configure another application using the wizard, you must first unconfigure the existing application monitoring configuration.

  Note: When you configure or unconfigure application monitoring, it does not affect the state of the application. The application runs unaffected on the virtual machine.

- After you have configured monitoring for an application using the wizard, you can configure monitoring for additional applications from VOM or the command line.
  For more information on how to use Veritas Cluster Server commands or VOM to configure additional applications, see the following technical note:
  http://www.symantec.com/docs/TECH159846
- If you clone a virtual machine on which you have configured application monitoring, you must reconfigure application monitoring on the cloned virtual machine.
- If a configured application fails, Symantec ApplicationHA attempts to restart the component on the virtual machine. If the component does not start, ApplicationHA reboots the operating system. If the application still does not start, ApplicationHA communicates with VCS to take corrective action. ApplicationHA then stops the other configured components in a predefined order. This prevents the other components from getting corrupted due to a machine reboot.
  Thus, a single failed component can bring down other healthy components running on the virtual machine. You must take this behavior into consideration while configuring application monitoring on a virtual machine.

Before configuring application monitoring

Note the following prerequisites before configuring application monitoring on a virtual machine:

- Verify that you have installed the Veritas Operations Manager (VOM) Management Server and the VOM add-on for ApplicationHA Management on the VOM Management Server.

You can also perform the application monitoring operations directly from a browser window by using the following URL:

```
https://<virtualmachineNameorIPaddress>:5634/vcs/admin/
application_health.html?priv=ADMIN
```

- Verify that you have installed Symantec ApplicationHA on the required virtual machines.
  Refer to the *Symantec ApplicationHA Installation Guide* for instructions.
- Verify that the logged-on user has administrative privileges on the virtual machine where you wish to configure application monitoring.
- Verify that you have appropriate user privileges on Veritas Operations Manager.
- Verify that the virtual machines on which you want to monitor applications are added as managed hosts to VOM.
- Verify that the ApplicationHA Application Monitoring Configuration Wizard has administrative credentials on the virtual machine for logging on, configuring, and administering applications on the virtual machine.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Symantec ApplicationHA installer, wizard, services, and VOM.
  Refer to the *Symantec ApplicationHA Installation Guide* for a list of ports and services used.

Chapter 4

Configuring VCS support for ApplicationHA

This chapter includes the following topics:

About VCS support for ApplicationHA

If you install Veritas Cluster Server (VCS) in the virtualization infrastructure layer and enable VCS to support ApplicationHA, you can administer application-aware monitoring of virtual machines by using VCS.

As part of application-aware monitoring of virtual machines, VCS performs the following actions, based on the application health status determined by ApplicationHA:

1. VCS restarts virtual machines if an application that is configured for monitoring with ApplicationHA faults. VCS restarts the virtual machine only if ApplicationHA is unable to bring the application back online by restarting

the application or by initiating a graceful internal restart of the virtual machine (soft reboot).

2. VCS fails over a virtual machine to another node in the same VCS cluster. VCS fails over the virtual machine only if the virtual machine restart fails to bring a faulted application online. For successful failover, both the source and destination nodes must fulfill certain requirements of storage space and network connectivity. For more information on permissions and ports and firewall settings for application-aware monitoring of virtual machines, refer the *Symantec ApplicationHA Installation Guide*.
3. VCS supports the live migration of virtual machines to another node in the same VCS cluster. VCS does not initiate any fault-management actions during such migration.

Enabling VCS support for ApplicationHA

To leverage the clustering capabilities of Veritas Cluster Server (VCS) to support ApplicationHA, you must run the enable_applicationha script on each VCS node (physical host).

The script sets up a private network between the VCS node and each virtual machine. Using this private network, ApplicationHA communicates application fault status to the VCS cluster.

You can also use the script to activate the auto-registration feature of a VCS cluster. For more information on the auto-registration feature:

See "About auto-registration of virtual machines with VCS node" on page 31.

Note: See " How ApplicationHA is deployed in the KVM environment" on page 12.. Also, to understand how VCS is deployed in the KVM environment, you must read the *Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide for Linux*

You can use one of the following methods to configure VCS support for ApplicationHA.

- Configuring VCS support for ApplicationHA
  See "Configuring VCS support for ApplicationHA using custom values" on page 32.
- Autoconfiguring VCS support for ApplicationHA
  See "Configuring VCS support for ApplicationHA using default values" on page 33.

- Configuring VCS support for ApplicationHA using response files
  See "Configuring VCS support for ApplicationHA using a response file" on page 34.

About auto-registration of virtual machines with VCS node

In the KVM virtualization technology, the auto-registration feature automatically sets up communication between the VCS node (physical host) and all the virtual machines where ApplicationHA is configured. Auto-registration also sets up a VCS resource (KVMGuest resource) for VCS to monitor the virtual machines. These steps enable VCS to take application-aware administrative actions such as start, stop, or fail over virtual machines.

Note: Enabling auto-registration is an optional step. Before you enable auto-registration, you must set up a private VLAN.

In the KVM environment, you can either specify the DHCP option or the VirtIO channel to enable the autoregistration feature:

The following table lists the comparitive advantages of the two options.

DHCP	VirtIO
virtual machine reboot not required for auto-registration	virtual machine reboot required for auto-registration
DHCP server must be running on physical host	No new infrastructure is not required
May not work if the guest name of the virtual machine is different from the host name	Works even if the guest name of the virtual machine is different from the host name
Supports only 254 virtual machines per physical host	Supports 254 X 254 virtual machines per physical host/
The enable_appHA script creates a resource for VCS to monitor the DHCP process .	Such monitoring is not required in the VirtIO mode.

Configuring VCS support for ApplicationHA using custom values

To manually configure VCS to support ApplicationHA:

1. Navigate to the following location and run the enable_applicationha script:

   `/opt/VRTSvcs/bin/utils`

2. When prompted, specify the value of some of the following private VLAN configuration parameters. Some of the following parameters may automatically be set to default values.

   The following table lists each parameter and its description:

Private VLAN configuration parameter	Description
Private VLAN ID	Specifies a unique identifier to distinguish private communication between the virtual machines and the VCS node (physical host). This identifier ensures that communication over the designated private network neither interferes with, nor is visible from the public network connecting the virtual machines and the VCS node. You can specify a private VLAN ID between 2 and 4094. The default value is 123.
Physical NIC on the physical host/frame	Specifies the NIC that the private VLAN uses for communication between the VCS node (physical host) and the virtual machines. You can specify even a public NIC on the physical host. You cannot specify a NIC that is used as: ■ LLT heartbeat interface ■ Loopback interface ■ Bridge interface
Virtual machine	Specifies the virtual machines that must be added to the private VLAN. The script lists all virtual machines that are not added to the VLAN. You can specify all or some of the listed virtual machines. Once you specify a virtual machine, the script creates a virtual interface for the virtual machine to communicate over the private VLAN.

Private VLAN configuration parameter	Description
Network	Specifies the network that the virtual interfaces must use for the private VLAN. You must specify the network in the following format: X.Y.0.0 Where the values of a and b lie between 1and 254. The network you specify must not be already in use by the physical host/frame. The default network is 192.168.0.0

3 At the appropriate prompt, accept the option to enable autoregistration.

Autoregistration mode	Description
DHCP	The physical host acts as a DHCP server and allots IP addresses to the virtual machines connected to the private network. Each virtual machine configured for application monitoring with ApplicationHA requests the DHCP server for such an IP address.
VirtIO	This is the native communication channel in the KVM virtualization environment. You must specify this option in case DHCP is not available. If you select VirtIO mode, then you must restart the virtual machine for the configuration to take effect.

4 To troubleshoot the configuration, see the following log file:

```
/var/VRTSvcs/log/applicationha_utils.log
```

Configuring VCS support for ApplicationHA using default values

When you configure VCS support for ApplicationHA, if you want to set default values for all configuration parameters, use the autoconfigure option of the `enable_applicationha` script.

In some cases, the script may prompt you to specify values for certain attributes.

Note: In case, the VCS cluster is secure, you must manually copy the credential file `/var/VRTSvcs/vcsauth/data/ApplicationVM.cred` to the `/var/tmp/` directory on all configured virtual machines. For security reasons, ensure that only the superuser has permissions to access this file.

To autoconfigure VCS support for ApplicationHA

1 Navigate to the following location:

```
/opt/VRTSvcs/bin/utils/
```

2 Run the following command:

```
# enable_applicationha -autoconfigure
```

Configuring VCS support for ApplicationHA using a response file

The `enable_applicationha` script supports response files. If you want to use the script as a non-interactive activity, you can specify custom values for the script parameters in a response file.

To use a response file to enable VCS support for ApplicationHA:

1 Create a response file by setting variables.

For more information:

See "Response file variables to enable VCS support for ApplicationHA" on page 35.

See "Sample response file for configuring VCS support for ApplicationHA" on page 37.

2 Run the following command:

```
# enable_applicationha -responsefile file
```

Where *file* is the name of the file you created in step 1.

Response file variables to enable VCS support for ApplicationHA

Table 4-1 Infrastructure variables

Parameter	Description
CFG{SYSTEMS}	Specifies the VCS node (physical host) where the script must run. You can specify only a local node.
CFG{UPGRADE}	Specifies whether the script must enable VCS support for ApplicationHA. To enable VCS support, you must set the value to 1.
CFG{VLANID}	Specifies a unique identifier for the private network between the VCS node (physical host) and the virtual machine.
CFG{DHCPCONFIGFILE}	Specifies the DHCP server configuration file on the VCS node (physical host) if anything other than default.
CFG{USE_CURRENT_VLAN}	Specifies whether an existing VLAN on the VCS node may be used for the private VLAN
CFG{ NIC}	Specifies the NIC to be used for the private VLAN
CFG{NETWORK}	Specifies the network name that is required when using the private VLAN
CFG{USE_CURRENT_NETWORK}	Specifies whether a network currently configured over the VCS node (physical host) must be used to create the private VLAN with virtual machines.
CFG{ AUTOREGISTRATION}	Specifies whether the autoregistration option must be used (that is, all specified virtual machines must be configured with default values), while enabling VCS support for ApplicationHA.
CFG{AUTOREGISTRATION _OPTION}	Specifies the mode of autoregistration. Set the value to 1 to specify the DHCP mode of autoregistration. Set the value to 2 to specify VirtIO mode of autoregistration. For more information on the VirtIO and DHCP modes, See "About auto-registration of virtual machines with VCS node" on page 31.

Table 4-1 Infrastructure variables *(continued)*

Parameter	Description
CFG{DHCPCONFIGFILE}	Specifies the location of the DHCP configuration file. Use this parameter in your response file if you specified DHCP mode of autoregistration. The default location of the DHCP file is: `/etc/dhcp/dhcpd.conf`
CFG{ RESET}	Rolls back the current ApplicationHA network configuration on the VCS node (physical host). Any virtual machines configured earlier are also removed from the private VLAN.

Table 4-2 Virtual Machine variables

Parameter	Description
CFG{ACTIVATE_VLAN_ALL}	Specifies if all virtual machines must be added to the private network. To add all virtual machines to the private network, set this parameter value to 1.
CFG{ACTIVATE_VLAN_GUESTS}	Specifies the virtual machines that must be added to the private network. You can specify multiple virtual machines by using a list in which names of the virtual machines are separated by a space " ". **Note:** Do not use this parameter in your response file if you set the ACTIVATE_VLAN_ALL parameter to 1.
CFG{REMOVE_VM_ALL}	Specifies if all virtual machines must be removed from the private VLAN. To remove the virtual machines from the private VLAN, set the value of this parameter to 1.
CFG{REMOVE_VM_GUESTS}	Specifies the virtual machines that must be removed from the private VLAN. You can specify multiple virtual machines by using a list in which names of the virtual machines are separated by a space " ".

Sample response file for configuring VCS support for ApplicationHA

Review the response file variables and their definitions and then create a response file.

See "Response file variables to enable VCS support for ApplicationHA" on page 35.

Following is a sample response file:

```
our %CFG;
$CFG{SYSTEMS}="appvcs01";
$CFG{VLANID}="123";
$CFG{AUTOREGISTRATION}="1";
$CFG{ACTIVATE_VLAN_ALL}="1";
$CFG{NIC}="eth0";
$CFG{UPGRADE}="1";
$CFG{NETWORK}="192.168.0.0";
$CFG{AUTOREGISTRATION_OPTION}="1";
```

Configuring VCS settings for ApplicationHA

Configuring Veritas Cluster Server (VCS) settings enables VCS to take specified actions, such as restarting the virtual machine, if the virtual machine does not send an application heartbeat within a configurable time interval.

VCS monitors the virtual machine through the KVM guest agent. The default monitoring frequency for the virtual machine is 60 seconds. In case of a heartbeat failure, the default action is to restart the virtual machine. You can configure the number of such attempts to restart. For more information, see the descriptions of the RestartLimit and ToleranceLimit attributes in the table below.

To configure VCS settings

Note: The values that you set in this procedure apply to all virtual machines running on a physical host.

---

1. In the Veritas Operations Managemer (VOM) Console, click **Manage** > **Server** > **Hosts**.
2. In the right pane, click the hostname/IP address of the VCS node where you want to configure monitoring settings.
3. In the **Service Groups** tab, click the name of the service group that contains the KVM resource.

4 To change the restart settings of the virtual machine, perform the following steps:

a. In the **Resources** tab, click the resource name of the virtual machine resource.

b. In the **All attributes** tab, right-click the following attribute, and from the context-menu select **Edit Attribute**:

Attribute	Description
RestartLimit	Specifies the number of times that VCS attempts to restart the virtual machine, before declaring it as faulted, and taking the next configured corrective action .

5 To change the monitoring settings of the virtual machine, perform the following steps:

a. In the **Resources** tab, click the resource **Type** (that is, KVMGuest)

b. In the All attributes **tab**, right-click the attribute that you want to edit, and from the context-menu, select Edit Attribute for changing settings for the following attributes:

Attribute	Description
MonitorInterval	Specifies the time in seconds after which the VCS virtualization agent checks for the state of the virtual machine, if the current state is Online.
OfflineMonitorInterval	Specifies the time in seconds after which the VCS virtualization agent checks for the state of the virtual machine, if the current state is Offline/Faulted.
ToleranceLimit	Specifies the number of monitor intervals for which VCS must wait for the heartbeat from a virtual machine to resume after a heartbeat failure, before declaring the virtual machine Offline.
OnlineRetryLimit	Specifies the number of times that the VCS virtualization agent tries to restart a virtual machine that fails to start. The default value is 0.
OnlineWaitLimit	Specifies the number of monitor intervals for which VCS waits for a virtual machine to come online after it is started, before reporting a fault.

Disabling VCS settings for ApplicationHA

You can disable Veritas Cluster Server (VCS) settings if you do not want VCS to restart a virtual machine in case of a heartbeat failure.

To disable VCS settings

1. In the Veritas Operations Managemer (VOM) Console, click **Manage** > **Servers** > **Hosts**.
2. Click the hostname/IP address of the VCS node where you want to unconfigure monitoring settings.
3. In the **Service Groups** tab, click the name of the service group that contains the KVM resource.
4. On the **All attributes** page, right-click the **ManageFaults** attribute , and from the context-menu select **Edit Attribute**
5. Set the value of the attribute to **None**.

Attribute	Description
ManageFaults	Enables VCS to restart a virtual machine, if the virtual machine fails. The default value 'ALL' enables VCS actions. You can set the value 'NONE' to disable VCS actions.

Chapter 5

Administering application monitoring

This chapter includes the following topics:

Administering application monitoring using the ApplicationHA tab

Symantec ApplicationHA provides an interface, the ApplicationHA tab, to configure and control application monitoring. The ApplicationHA tab is integrated with the Veritas Operations Manager.

Use the ApplicationHA tab to perform the following tasks:

- Configure and unconfigure application monitoring
- Start and stop configured applications
- Enable and disable application heartbeat
- Enter and exit maintenance mode

To view the ApplicationHA tab, launch the Veritas Operations Manager console, navigate to the virtual machine on which you want to monitor an application. For more information: See "Accessing the ApplicationHA tab" on page 22.

Note: You can also perform the application monitoring operations directly from a browser window using the following URL:
https://*<VMNameorIP>*:5634/vcs/admin/application_health.html?priv=ADMIN
where *<VMNameorIP>* is the virtual machinename or the IP address.

To configure or unconfigure application monitoring

Use the ApplicationHA tab to configure or unconfigure an application monitoring configuration from the virtual machine. This may be required in case you wish to re-create the configuration or configure another application using the wizard.

You can click the following links:

- Click **Configure Application Monitoring** to launch the Symantec ApplicationHA Configuration Wizard. Use the wizard to configure application monitoring.
- Click **Unconfigure Application Monitoring** to delete the application monitoring configuration from the virtual machine.
  Symantec ApplicationHA removes all the configured resources for the application and its services.
  Note that this does not uninstall Symantec ApplicationHA from the virtual machine. This only removes the configuration. The unconfigure option removes all the application monitoring configuration resources from the virtual machine.

To view the status of configured applications

Under the Component List tab of the VOM console, the Description box in the ApplicationHA displays the status of the configured application and the associated services.

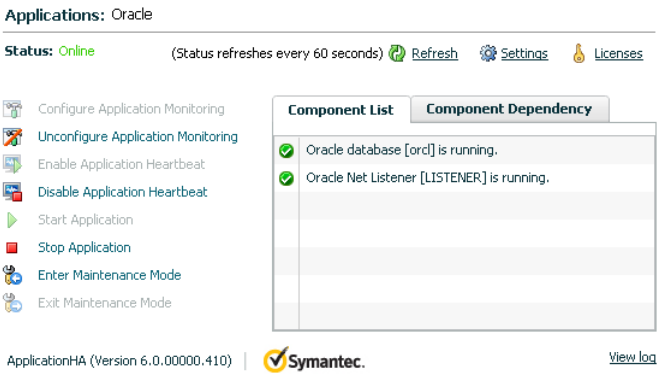

For example, if you have configured monitoring for Oracle, the Description displays the following information:

```
Oracle Net Listener [Listener] is running.
```

The Description box also displays the state of the configured application and its components. The following states are displayed:

online	Indicates that the services and processes are running on the virtual machine.
offline	Indicates that the services and processes are not running on the virtual machine.
partial	Indicates that either the services and processes are being started on the virtual machine or ApplicationHA was unable to start one or more of the configured services or processes.
faulted	Indicates that the configured services or components have unexpectedly stopped running

Click **Refresh** to see the most current status of the configured components. The status is refreshed every 60 seconds by default.

To view component dependency

ApplicationHA monitors a configured application for high availability by monitoring the status of its components. Inter-related components form a component group. The status of the application depends on the component groups.

The Component Dependency tab of the Veritas Operations Manager console illustrates this dependency between the application and its components

To access the Component Dependency tab perform the following step:

- In the ApplicationHA tab of the VOM client GUI, click **Component Dependency**. A component dependency graph appears. The graph indicates which component depends on which other component to be up and running.

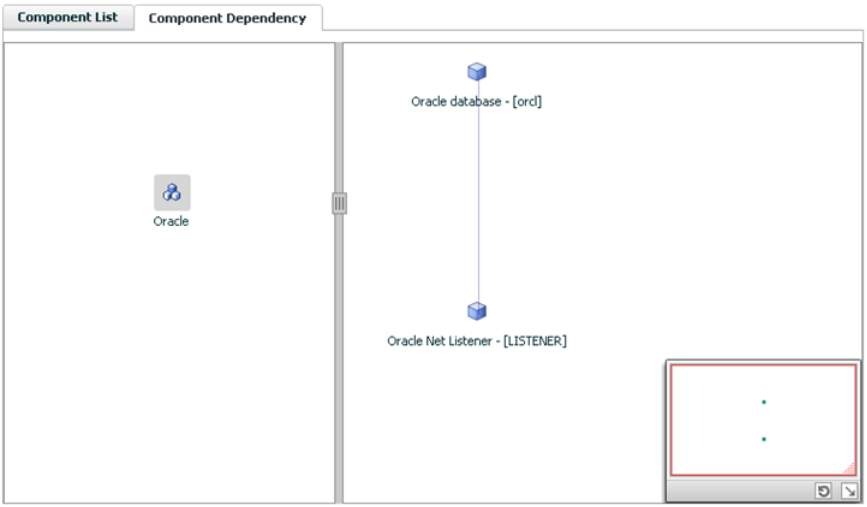

The above figure illustrates the component dependency for Oracle Database.

The left pane indicates the component groups, while the right pane indicates the components of the selected component group. A vertical line joining two components indicates that for the component at the higher level to be running, the component at the lower level must be running.

The track pad, at the left-bottom corner helps you navigate through complex component dependency graphs.

The Component Dependency graph also uses various icons to indicate component groups and components as well as their status. If you roll the mouse over a component, a tooltip highlights the name and the state of the component.

For example, in the above figure the Oracle component group is selected in the left pane and the components, Oracle Net Listener and Oracle database, appear in the right pane. For the Oracle database to be running, Oracle Net Listener must be up and running.

To start or stop applications

Use the following options on the ApplicationHA tab to control the status of the configured application and the associated components:

- Click **Start Application** to start a configured application.
  Symantec ApplicationHA attempts to start the configured application and its components in the required order. The configured resources are also brought online in a predefined order.

- Click **Stop Application** to stop a configured application that is running on the virtual machine.
  Symantec ApplicationHA begins to stop the configured application and its components gracefully. The configured resources are also taken offline in the pre-defined order.

To enable or disable application heartbeat

The VCS KVMGuest agent uses the heartbeat information that VCS captures as a proxy for virtual machine availability. This allows VCS to automatically restart individual virtual machines that have lost their ability to heartbeat either due to application failure or due to virtual machine failure.

The ApplicationHA tab allows you to control the application heartbeat on the virtual machines.

Use the following options on the ApplicationHA tab to control the configured application heartbeat:

- Click **Enable Application Heartbeat** to enable the heartbeat communication between the configured applications running on the virtual machine and VCS. The application heartbeat is enabled by default when an application is configured for monitoring.
- Click **Disable Application Heartbeat** to disable the heartbeat communication between the configured applications running on the virtual machine and VCS. If you disable application heartbeat, VCS does not restart the virtual machine when application in the virtual machine fails.

To suspend or resume application monitoring

After configuring application monitoring you may want to perform routine maintenance tasks on those applications. These tasks may or may not involve stopping the application but may temporarily affect the state of the applications and its dependent components. If there is any change to the application status, Symantec ApplicationHA may try to restore the application state. This may potentially affect the maintenance tasks that you intend to perform on those applications.

If stopping the application is not an option, you can suspend application monitoring and create a window for performing such maintenance tasks. When application monitoring is suspended, ApplicationHA freezes the application configuration, disables the application heartbeat, and stops sending the heartbeat to VCS.

The ApplicationHA tab provides the following options:

- Click **Enter Maintenance Mode** to suspend the application monitoring for the applications that are configured on the virtual machine. During the time the monitoring is suspended, Symantec ApplicationHA does not monitor the state of the application and its dependent components. The ApplicationHA tab does not display the current status of the application. If there is any failure in the application or its components, ApplicationHA takes no action.
- Click **Exit Maintenance Mode** to resume the application monitoring for the applications configured on the virtual machine. You may have to click the **Refresh** link in the ApplicationHA tab to see the current status of the application.

  When application monitoring is restarted from a suspended state, ApplicationHA does not enable the application heartbeat. Click **Enable Application Heartbeat** to enable it.

  If you have made changes to the application that is being monitored, then those changes may not reflect in the application monitoring configuration. In such cases, you may have to unconfigure and reconfigure the application monitoring.

Administering application monitoring settings

The ApplicationHA view provides a set of options that you can use to control the way Symantec ApplicationHA handles application monitoring, application and dependent component faults, and application recovery on the virtual machine. The view also provides a set of options that you can use to configure ApplicationHA to restart the virtual machine. These configuration settings are applicable on a per virtual machine basis. The settings apply to all the applications that Symantec ApplicationHA monitors on the virtual machine.

The following settings are available:

- App.StartStopTimeout

  When you click the **Start Application** or **Stop Application** links in the ApplicationHA view, Symantec ApplicationHA initiates an orderly start or stop of the application and its dependent components. This option defines the number of seconds Symantec ApplicationHA must wait for the application to start or stop. If the application does not respond in the stipulated time, an error is displayed in the ApplicationHA view.

  A delay in the application response does not indicate that the application or its dependent component has faulted. Parameters such as workload, system performance, and network bandwidth may affect the application response. Symantec ApplicationHA continues to wait for the application response even after the timeout interval is over. If the application fails to start or stop,

ApplicationHA takes the necessary action depending on the other configuration settings.

AppStartStopTimeout value can vary between 0 and 600. The default is 30 seconds.

- App.RestartAttempts

  This option defines the number of times Symantec ApplicationHA should try to restart a failed application or its dependent component. If an application fails to start in the specified number of attempts, Symantec ApplicationHA stops the application heartbeat and communicates the fault to VCS.

  AppRestartAttempts value can vary between 1 and 6. The default is 1.

- App.ShutdownGraceTime

  This option defines the number of seconds Symantec ApplicationHA should wait before communicating the application fault to VCS.

  If a configured application or its dependent component fails, Symantec ApplicationHA tries to restart the component for the configured number of times. If the component fails to start, Symantec ApplicationHA stops the application heartbeat and communicates the fault to VCS. VCS may then restart the virtual machine depending on the configuration settings.

  An abrupt shutdown may affect the other healthy application components running on the machine. If those components require more time to stop, Symantec ApplicationHA may not be able to stop them gracefully in time before the reboot is initiated. For such cases, you can use AppShutdownGraceTime to delay the virtual machine reboot so that Symantec ApplicationHA stops all the application components gracefully.

  When an application fails to start, Symantec ApplicationHA initiates a graceful shutdown of all the healthy applications being monitored on the virtual machine and waits for time specified in this option. A virtual machine reboot takes place only after all the application components are shut down gracefully or at the end of the grace time, whichever is earlier.

  This setting is applicable to the heartbeat service group that is created when you configure application monitoring using the Symantec ApplicationHA Configuration Wizard. Internally, it sets the DelayBeforeAppFault attribute of the Heartbeat agent resource (VCSAppMonHBRes) in the configuration.

  AppShutDownGraceTime value can vary between 0 and 600. The default is 300 seconds.

- VM.GracefulRebootPolicy

  Use this option to enable or disable ApplicationHA-initiated virtual machine restart policy. This option defines whether or not ApplicationHA restarts the virtual machine in response to application and component failures. When a configured application or component fails, ApplicationHA attempts to restart

the failed components. If the component fails to start, ApplicationHA then takes the next corrective action.
If this policy is disabled, and an application or component fails, then ApplicationHA stops sending the heartbeat to VCS. As a result VCS can then restart the virtual machine.
If this policy is enabled, ApplicationHA itself invokes a native operating system command to restart the virtual machine.
VM.GracefulRebootPolicy value can be Enabled (1) or Disabled (0). The default value is Disabled.
ApplicationHA uses the following command to initiate a soft reboot:

```
# /sbin/reboot
```

To display a broadcast message that tells users about the soft-internal reboot in progress, perform the following steps:
1. In the Veritas Operations Managemer (VOM) console, click **Manage** > **Server** > **Hosts**.
2. In the right pane, click the hostname/IP address of the virtual machine where you want to update the soft reboot command.
3. In the **Service Groups** tab, click on the infrastructure service group name, that is VCSAppMonHBSG.
4. In the **Resources** tab, click the heartbeat resource name, that is VCSAppMonHBRes.
5. In the All attributes tab, right-click the following attribute, and from the context-menu select **Edit Attribute**:

Attribute	Description
VMGracefulRebootCmd	Specifies the command that ApplicationHA to perform soft reboot of the virtual machine.

You can set the following value for the attribute to display a custom message message “ApplicationHA initiated soft reboot” on the console:

```
/sbin/shutdown -r now ApplicationHA initiated soft reboot
```

- VM.GracefulRebootAttempts
  This option defines the number of times ApplicationHA attempts to restart the virtual machine gracefully if the configured application or component becomes unresponsive. The number of restart attempts is time bound and is defined by the option VM.GracefulRebootTimeSpan. The restart attempts count is reset after the reboot time span elapses.
  For example, if the reboot attempts value is 4, the time span value is 1 hour, and ApplicationHA has restarted the virtual machine once, then the restart attempt count is 3 (initial set value of 4 minus one reboot) for the remaining

period of the 1-hour interval. The restart attempts count is reset to 4 at the beginning of the next 1-hour span.
If the restart attempts are exhausted and the application or component fails within the reboot time span again, ApplicationHA stops the application heartbeat and communicates the fault to VCS. Depending on the configuration, VCS may then restart or fail over the virtual machine.
VM.GracefulRebootAttempts value can vary between 1 and 10. The default value is 1.

- VM.GracefulRebootTimeSpan
  This option defines the time interval, in hours, during which ApplicationHA can gracefully restart the virtual machine for the number of times defined by the option VM.GracefulRebootAttempts.
  VM.GracefulRebootTimeSpan value can vary between 1 and 24. The default value is 1 hour.

Note: These attribute values are not affected due to a hard restart by VCS. The configuration remains in effect even after VCS reboots the virtual machine.

To modify the application monitoring configuration settings

1. Launch the Veritas Operations Manager (VOM) console, and navigate to the virtual machine where you have configured application monitoring.

   For more information on navigating the VOM console, See "Accessing the ApplicationHA tab" on page 22.

2. Select the **ApplicationHA** tab and then click the **Settings** link to display the Settings dialog box.

3. Specify the values for the available options displayed in the Settings box and then click **OK**.

   The specified values are updated in the configuration and they take effect immediately.

About ApplicationHA-initiated virtual machine restarts

When you configure application monitoring, ApplicationHA uses a heartbeat to communicate the application status to VCS in the virtualization infrastructure layer. If the application or its component fails, ApplicationHA attempts to restart it. If the application does not start, depending on the configuration, ApplicationHA initiates a soft reboot of the virtual machine. If the application does not come

online after the soft reboot, ApplicationHA stops sending the heartbeat to VCS. Depending on the configuration, VCS in the infrastructure layer then performs a hard restart of the virtual machine.

A hard restart has various implications and may not be the desired solution at all times. It may prove to be counter productive in several cases, including the following:

- The virtual machine itself is running fine but the application is unable to get the required resources
- The other applications and tools running on the machine may either hang or take longer time to recover or restart after an abrupt shut down
- A hard restart can be potentially disruptive if there are multiple critical applications running on the virtual machine

A soft reboot or an operating system driven restart is graceful and allows for a more orderly shutdown of applications and tools running on the machine and can help reduce potential disruption to critical applications. ApplicationHA provides this layer of application control wherein you can configure ApplicationHA itself to restart the virtual machine using native operating system commands. Note that the soft reboot of the virtual machine by ApplicationHA is disabled by default.

ApplicationHA provides attributes that you can use to configure ApplicationHA to restart the virtual machine.

See "Administering application monitoring settings" on page 46.

Does ApplicationHA-initiated reboot affect VCS HA?

ApplicationHA-initiated reboot works independently of high availability provided by VCS installed in the virtualization infrastructure layer. It is not intended to replace VCS in the infrastructure layer. It is useful in cases where there is a need to first bring down other healthy applications and dependent components before actually restarting the virtual machine.

ApplicationHA-initiated reboot complements VCS in the infrastructure layer by offering an additional layer of control that helps in building customized application management and recovery plans in virtualization environments. ApplicationHA-initiated reboot can co-exist with VCS in the infrastructure layer. You can configure both ApplicationHA-initiated reboot and VCS in the infrastructure layer as part of your recovery plan. ApplicationHA-initiated reboot can act as the first line of action against application failures. If a graceful restart does not resolve the application failures, then depending on the reboot configuration settings, ApplicationHA stops the application heartbeat and VCS in the infrastructure layer then takes control of the virtual machine.

Chapter 6

Administering VCS support for ApplicationHA

This chapter includes the following topics:

About administering VCS support for ApplicationHA

When you enable VCS support for ApplicationHA, VCS extends its ability to externally reboot a faulted virtual machine or fail over a faulted machine to another node in the VCS cluster.

When you enable VCS support for ApplicationHA, virtual machines that are configured for application monitoring with ApplicationHA, are registered with VCS. As part of the registration, a VCS resource is created on each virtual machine. If a resource already exists (for application-status agnostic monitoring of the virtual machine), VCS uses the same resource to execute ApplicationHA tasks.

You can perform the following administrative tasks using VCS to monitor configured virtual machines:

- Configuring a new virtual machine for application-aware monitoring

See “Configuring a new virtual machine for application-aware monitoring ” on page 52.

- Unconfiguring application-aware monitoring of a virtual machine
  See “Unconfiguring application-aware monitoring of a virtual machine” on page 53.
- Viewing configuration details of a virtual machine
  See “Viewing configuration details of a virtual machine” on page 54.
- Viewing connection details of a virtual machine
  See “Viewing connection details of a virtual machine” on page 55.
- Putting a physical host into maintenance mode
  See “Putting a physical host into maintenance mode” on page 56.

Configuring a new virtual machine for application-aware monitoring

If you add a new virtual machine to an existing ApplicationHA configuration, the new virtual machine is not automatically configured for application-aware monitoring by VCS. You must add the new virtual machine to the private network (VLAN) between the configured virtual machines and physical host. You can do so by running the `enable_applicationha` script with the following options.

Note: You must perform the following steps on each VCS node (physical host) in the cluster where you added new virtual machines for application-aware monitoring.

To add a new virtual machine to the private VLAN

1 If you want to configure all the newly created virtual machines for application-aware monitoring, run the following command:

```
# enable_applicationha -addvm ALL
```

The procedure for configuring all newly-created virtual machines is complete.

If you want to configure only some new virtual machines to the cluster, run the following command:

```
# enable_applicationha -addvm
```

2 Review the displayed list of new virtual machines and specify the names of only those virtual machines that you want to configure for application-aware monitoring, by using the following command:

```
# enable_applicationha -addvm vm1 vm2 ...
```

Where *vm1*, *vm2* are the names of the virtual machines that you want to add to the VCS cluster.

When you perform the above step/s, the virtual machines that are configured for monitoring with ApplicationHA, are automatically registered with the private VLAN when you run the above commands.

Unconfiguring application-aware monitoring of a virtual machine

If you want to perform some maintenance activities on a certain virtual machine, you may want to suspend application-aware monitoring of that virtual machine by VCS. This prevents VCS from taking any fault-management steps during the maintenance activity.

You can achieve this by removing the virtual machine from the private VLAN between configured virtual machines and the VCS node (physical host).

Note: You must perform the following steps on each VCS node (physical host) in the cluster where you want to suspend virtual machines application-aware monitoring of virtual machines by VCS.

To remove a virtual machine from the private VLAN

1 If you want to remove all registered virtual machines from the private VLAN, run the following command:

```
# enable_applicationha -delvm ALL
```

The procedure to remove all virtual machines from the private VLAN is complete. If you want to remove certain virtual machines from the VCS cluster, run the following command:

```
# enable_applicationha -delvm
```

2 Review the displayed list of registered virtual machines in the private VLAN, and specify those virtual machines that you want to remove from the private VLAN:

```
# enable_applicationha -delvm vm1 vm2
```

Where *vm1*, *vm2* are the names of the virtual machines where you want to suspend application-aware monitoring by VCS.

Once you remove a virtual machine or multiple virtual machines from the private VLAN by using the above steps, they do not automatically register with the private VLAN as per the autoregistration feature.

If you want to restart application-aware monitoring on such virtual machines, you must repeat the steps described in the following topic:

See "Configuring a new virtual machine for application-aware monitoring " on page 52.

Viewing configuration details of a virtual machine

You can view the configuration details, including infrastructure parameters such as NIC and VLAN ID, of various virtual machines associated with a VCS node (physical host).

- To view configuration details of all virtual machines associated with a VCS node (physical host), run the following command:

```
# enable_applicationha -status
```

You can also view the configuration details of a specific virtual machine from the command line. The output displays status in terms of "Registered" or "Not Registered". The 'Registered' status implies that the infrastructure for

application-aware monitoring of that virtual machine is created. However, the virtual machine may not be connected to the VCS node (physical host).

- To view the registration details of selected virtual machines, run the following command:

  ```
  # enable_applicationha -status vm1 vm2 ...
  ```

  Where vm1, vm2 are names of virtual machines of which you want to view configuration details.

Viewing connection details of a virtual machine

You can view the connection details of a virtual machine from the VCS node (physical host).

- To view the connection state of a virtual machine, run the following command:

  ```
  # hasys -value VirtGuest ConnectionState
  ```

  Where *VirtGuest* is the host name of the virtual machine.
  The host name of the virtual machine is reflected in the "CEInfo" attribute of the KVMGuest resource corresponding to the virtual machine. To get this information, use the command:

  ```
  # hares -value virt_name CEInfo
  ```

  Where *virt_name* is the name of virtual machine. In the output, the name that corresponds to the key CESystem, is the host name of the virtual machine.
  The following table lists each connection state value and its description.

Connection state	Description
CONNECTED	The virtual machine is connected to the physical host.
NOT RESPONDING	The virtual machine has stopped responding to the VCS cluster. VCS does not take any fault-management steps.
DISCONNECTED	The virtual machine is not connected to the VCS cluster. Depending on the previous state of the virtual machine, VCS takes fault-management steps.

Connection state	Description
DISABLED	The connection between the virtual machine and the physical host is disabled. VCS does not perform any application-aware monitoring of the virtual machine.

- To determine the VCS node (physical host) to which a virtual machine is connected, run the following command. At any given time, a virtual machine is connected to only one VCS node:

  ```
  # hasys -value VirtGuest ControllerNode
  ```

  Where *VirtGuest* is the host name of the virtual machine

Putting a physical host into maintenance mode

You can use the `enable_applicationha` script to put a VCS node (physical host), that is configured to support ApplicationHA, into maintenance mode.

You can do so by using the following command:

```
# enable_applicationha -reset
```

This configuration step resets the ApplicationHA configuration on the VCS node where you run the script. You can put all the nodes in a VCS cluster, one by one, into maintenance mode.

When you put the last node into maintenance mode, global ApplicationHA settings such as IP addresses and DHCP resources are also deleted across the VCS cluster.

ApplicationHA, however, continues to monitor applications on the configured virtual machines

Chapter 7

Managing Symantec ApplicationHA licenses

This chapter includes the following topics:

About managing ApplicationHA licenses

When the embedded, two-month, evaluation license key expires, you may want to add a permanent license key.

You can add or view the license key from any virtual machine that has ApplicationHA guest components installed. You can use one of the following methods to manage the licenses:

- From the command line, run the following commands:
  To view an existing license:
  `/opt/VRTS/bin/vxlicrep`
  To install a new license:
  `/opt/VRTS/bin/vxlicinst`
- When you run the CPI installer to install or upgrade ApplicationHA, you can specify a new license key.
- Connect to the Veritas Operations Manager console and select the virtual machine for which you want to update the licenses. Select the **ApplicationHA** tab and click **Licenses**. Use this path to manage licenses for the local virtual machine.
  See "Managing ApplicationHA licenses through ApplicationHA tab" on page 58.

Managing ApplicationHA licenses through ApplicationHA tab

Perform the following steps to manage ApplicationHA licenses through the ApplicationHA tab.

To manage the ApplicationHA licenses

1. Connect to the Veritas Operations Manager.
2. In the Veritas Operations Manager console, click **Manage** > **Servers** > **Hosts**.
3. In the left pane, in the **License** list box, select the **ApplicationHA** check box.
4. In the right pane, click the virtual machine where you want to perform administrative actions.
5. Click the **ApplicationHA** tab and then click **Licenses**.
6. On the License Management panel, enter the new license key in the **Enter license key** text box and then click **Add**.
7. Click **Close**.

Appendix A

Troubleshooting Symantec ApplicationHA configuration

This appendix includes the following topics:

ApplicationHA view logging

The ApplicationHA view generates log files that are appended by letters. The log files are segregated based on operations and configuration settings, as follows:

- Operations and wizard logging
  ApplicationHA logs operations logs include the Symantec ApplicationHA Configuration Wizard logs and logs related to the various operations performed from the ApplicationHA view.
  Operations logs are located at: `/var/VRTSvcs/log`
  For example: `/var/VRTSvcs/log/AppControlOperations_A.log`
  The Symantec ApplicationHA Configuration Wizard also maintains in-memory logs that are available only during the time the wizard is running. These logs are maintained on a per session basis. The in-memory logs are purged after the wizard is closed. These logs are not stored in any file or directory.
- Configuration settings logging
  Application monitoring configuration settings related changes are logged separately and are available at:
  `/var/VRTSvcs/log`

For example: `/var/VRTSvcs/log/AppControlSettings_A.log`

These settings are accessible from the Settings link on the ApplicationHA view.

- ApplicationHA view logging
  The ApplicationHA view also maintains in-memory logs of the operations performed from the view. These logs are available only until the time the logs window is open. To view the current logs, click the **View Logs** link available on the right hand side in the ApplicationHA view. A window appears within the view. This window displays the details of the operations performed.

Symantec ApplicationHA tab does not display the application monitoring status

The Symantec ApplicationHA tab in the Veritas Operations Manager (VOM) console may either display a HTTP 404 Not Found error or may not show the application health status at all.

Verify the following conditions and then refresh the ApplicationHA tab in the VOM console:

- Verify that the ApplicationHA add-on is configured in VOM.
- Verify that the Veritas Storage Foundation Messaging Service (xprtld process) is running on the virtual machine.
  To verify xprtld process is running, run the following command:

  ```
  # /usr/bin/svcs svc:/system/xprtld:default
  ```

  If the process has stopped, run the following command:

  ```
  # svcadm enable svc:/system/xprtld:default
  ```

- Verify that ports 14152, 14153, and 5634 are not blocked by a firewall.
- Log out of VOM and then login again. Then, verify that the Symantec ApplicationHA plugin is installed and enabled.
  If the problem persists, perform the following steps in the VOM console:
  Click **Settings** > **Host Management**.
  Right-click the host where you want to view the monitoring status, and from the context menu click **Refresh Host(s)**.

Symantec ApplicationHA tab displays a "Failed to retrieve status" popup message

The Symantec ApplicationHA tab in the Veritas Operations Manager (VOM) console may display the following error in a popup window:

```
Failed to retrieve status.

Please ensure the machine is powered on and required services are running.
```

This error may occur if you reinstall or repair Symantec ApplicationHA Console in your application monitoring environment.

Perform the following actions:

- Verify that the virtual machine is powered on and accessible over the network.
- Verify that the Veritas Storage Foundation Messaging Service (xprtld) is running on the virtual machine.
- Close the ApplicationHA tab and open it again.
  In the VOM, click another virtual machine, then click the original virtual machine again and then select the ApplicationHA tab, or exit the VOM and launch it again.
  The ApplicationHA view then displays the status of the configured applications on the virtual machine.

Symantec ApplicationHA Configuration Wizard displays blank

The Symantec ApplicationHA Configuration Wizard may fail to display the wizard panels. The window may appear blank.

Workaround

You must re-launch the wizard.

VCS does not detect configured virtual machine

VCS does not detect a configured virtual machine because ApplicationHA installed on the virtual machine is not able to communicate with underlying VCS in the virtualization infrastructure layer.

This occurs when the following conditions simultaneously occur:

- You have installed ApplicationHA in a KVM virtualization environment.

- SELinux is enabled on the virtual machine.

(2483903)

Workaround

Disable SELinux on the virtual machine. For more information on disabling SELinux, see Red Hat Linux Enterprise documentation.

A virtual machine is unable to connect with VCS cluster

In a KVM environment, when a virtual machine exits maintenance mode, it is unable to re-establish communication with the VCS cluster that supports the virtual machine. (2566645)

Workaround

After the virtual machine exits maintenance mode, enable heartbeat for the virtual machine from the ApplicationHA tab in Veritas Operations Manager. For more information:

See "To enable or disable application heartbeat" on page 45.

A virtual machine loses communication over physical host

If you reset ApplicationHA configuration on the physical host, when a virtual machine is online, then the virtual machine is unable to communicate over VLAN, once the reset is complete. An internal reboot of the machine does not restore the communication channel. (2490026)

Workaround

Ensure the following:

- Thevirtual machine is stopped before you reset the ApplicationHA configuration.
- Restart the virtual machine from the physical host to restore the communication between virtual machine and the physical host.

ApplicationHA-initiated reboot does not broadcast any message on console

ApplicationHA-initiated reboot uses the default reboot command to reboot the virtual machine. If you want to alert a user with a broadcast message on the virtual machine console at the time of reboot, you can use the VM.GracefulRebootPolicy attribute.

For more information

See "Administering application monitoring settings" on page 46.

(2586314)

KVMGuest group faults on all systems

When you configure an application or modify the configuration, the KVMGuest group faults on all physical hosts:

Workaround

This may be due to an application misconfiguration.

From the physical host, perform the following steps:

1. Clear the KVMGuest service group fault.
2. Start the virtual machine outside VCS control.
3. Modify the application monitoring configuration.
4. Unfreeze the service group from the physical host.

ApplicationHA fails to restart an application

A configured application fails to restart on a virtual machine.

Workaround

Verify the value of 'App.RestartAttempts' attribute from the **Settings** menu of the option in the ApplicationHA tab of Veritas Operations Manager:

1. Verify if the value is greater than 0.

2. Verify if the appliation was restarted by ApplicationHA very recently. The value of the App.RestartAttempts attribute is effective only if the application has been up and running for a configurable period of time.

Soft reboot of a virtual machine is not triggered

ApplicationHA fails to restart a virtual machine even if the configured application running on the virtual machine fails.

Workaround

In the Settings menu on the ApplicationHA tab of Veritas Operations Manager, ensure that the value of the VM.GracefulRebootPolicy attribute is set to 'Enabled'. Ensure that ApplicationHA has not exceeded the number of restart attempts specified in the VM.GracefulRebootAttempts attribute.

Hard reboot of a virtual machine is not triggered

When a configured application fails on a virtual machine, VCS fails to externally reboot the virtual machine.

Workaround

- Verify that VCS on the physical host is enabled for ApplicationHA support.
- Ensure that the virtual machine is added to the private VLAN of the physical host.
- Ensure that the state of the virtual machine in the physical host is 'Connected'. For more information on viewing the connection state of a virtual machine:
  See "Viewing connection details of a virtual machine" on page 55.
- Ensure that the value of the RestartLimit attribute is greater than 0 for the resource corresponding to the virtual machine.

VCS cannot fail over virtual machine

VCS is unable to fail over a virtual machine even if an application running on the virtual machine fails, and other fault-management steps also fail.

Workaround

- Verify if VCS is enabled for ApplicationHA support on all physical host.
- Verify if the faulted virtual machine is part of the private VLAN of the physical host.
- Verify if the connection state of the virtual machine with the physical host is 'Connected'.
  For more information, viewing the connection state of a virtual machine:
  See "Viewing connection details of a virtual machine" on page 55.

Unconfiguring monitoring does not restore default application monitoring settings

When you unconfigure application monitoring on a virtual machine, the factory settings for ApplicationHA attributes, such as App.RestartAttempts and VM.GracefulRebootPolicy, are not automatically restored. Instead, the values that you configured before unconfiguring application monitoring on that virtual machine, are restored.

Workaround

This is expected behavior. When you reconfigure application monitoring on the virtual machine, if you want to set ApplicationHA to default values, you must manually reset the default values from the ApplicationHA tab.

VCS may fail over an online virtual machine

If a virtual machine comes online, but the KVMGuest resource does not come online, VCS fails over the virtual machine.

The following error message appears in the VCS error log:

```
KVMGuest resource name is disconnected. Returning OFFLINE.
```

Workaround

1. Ensure that ApplicationHA is configured and able to start on the virtual machine after a reboot.
2. If ApplicationHA takes unusually long to start on the virtual machine, increase the OnlineWaitLimit for the KVMGuest resource type.

   For more information on setting this attribute, See "Configuring VCS settings for ApplicationHA" on page 37.
3. If you do not want to monitor the virtual machine with ApplicationHA, unconfigure application-aware monitoring on the virtual machine.

   For more information, See "Unconfiguring application-aware monitoring of a virtual machine" on page 53.

Unconfiguring application-aware monitoring on a virtual machine may fail

If you want to stop application-aware monitoring by VCS on a virtual machine, you can do so by running the following command on the virtual machine:

```
# enable_applicationha -delvm vm1
```

Where *vm1* is the host name of the virtual machine where you want to stop application-aware monitoring by VCS. In some cases, this command may fail. (2575354)

Workaround

1. Unconfigure application monitoring from the ApplicationHA tab on VOM console, and then retry.
2. If step 1 fails, run the following command on the virtual machine and then retry:

   ```
   # /opt/VRTSvcs/bin/utils/remove_ip
   ```

ApplicationHA does not restart a failed application

If a configured application fails, ApplicationHA must restart it as configured. In some cases, this may fail.

Workaround

- Ensure that the value of App.RestartAttempts attribute for the virtual machine is greater than 0. You can review the value from the **Settings** option in the ApplicationHA tab of the VOM console.
- Ensure that ApplicationHA did not recently restart the application. If ApplicationHA recently restarted the application, then there may be a certain lag in time before the full value of App.RestartAttempts is restored. During this lag, if a fault re-occurs, ApplicationHA may not restart the application.

User is unable to add a virtual machine to private VLAN for monitoring

To be able to monitor a virtual machine for high availability with VCS, you must add the new virtual machine to the private VLAN between virtual machines and the associated VCS node (physical host). In some cases, this task may fail.

Workaround

1 From the virtual machine configuration, verify if slot “1F” for the PCI device is already in use.

Run the following commands on the virtual machine:

```
# virsh edit VirtName
```

Where *VirtName* is the name of the virtual machine. If the device using this slot is not a network bridge interface with “vcsbr” as the bridge, then you must assign the device to a slot lesser than “1F”.

Perform this step, and then add the virtual machine to the private VLAN.

2 If the first step does not work, and the virtual machine is running, shut down the virtual machine, and then add it to the private VLAN.

User is unable to add Virtual IO channel control device to virtual machine

When you enable VCS support for ApplicationHA, if you specify the VirtIO mode of autoregistration, the script adds the Virtual IO channel controller device to the virtual machine. This step may sometimes fail.

Workaround

1 From the virtual machine configuration, verify if the PCI device slot “1e” is already in use.

```
# virsh edit VirtName
```

or

```
# virsh dumpxml VirtName
```

Where *VirtName* is the name of the virtual machine. If the ID is in use and the device using the slot is not the Virt IO controller, you must assign a slot lesser than “1e” to the device.

2 If the first step does not work and the virtual machine is running, shut down the virtual machine, and then configure the virtual machine for autoregistration.

Migration of virtual machine fails

After the live migration of a virtual machine, if the backend boot image of the virtual machine uses the CFS, the virtual machine may suspend on the destination physical host.

Workaround

Ensure that the virtual machine boot image on CFS has all file permissions.

'virsh' command hangs on the physical host

If you run the `virsh` command on the physical host, it may hang without displaying an output.

Workaround

If the libvrtd service is not running, use the following command to start the service.

```
# Service libvirtd restart
```

Network interface not visible inside a registered virtual machine

After you add a virtual machine to the private VLAN between the virtual machines and the associated VCS node (physical host), the network interface inside the virtual machine may not be visible.

When you add a virtual machine to the private VLAN, a network interface is hot-plugged into the virtual machine. The interface is visible if you run the following command on the virtual machine:

```
# ifconfig - a
```

On virtual machines that are running RHEL 5, the hot-plugging may fail and the interface may not be visible.

Workaround

1. Remove the virtual machine from the private VLAN between the virtual machines and the associated VCS node (physical host). For more information:

   See "Unconfiguring application-aware monitoring of a virtual machine" on page 53.

2. Use the following command to load the kernel module "acpiphp":

   ```
   # modprobe acpiphp
   ```

3. Add the virtual machine to the private VLAN. For more information:

   See "Configuring a new virtual machine for application-aware monitoring " on page 52.

4. If the interface is now visible on the virtual machine, add the “acpiphp” module into the startup sequence of the virtual machine.
5. If steps 1 through 4 fail, shut down the VirtMachine and then add it to the private VLAN.

A virtual machine intermittently displays “NOT RESPONDING” status

If you query the connection status of a virtual machine from the command line, the virtual machine intermittently displays the status “NOT RESPONDING”. This may occur if:

- The virtual machine is overloaded
- The physical host is overloaded

As a result, there is a delay in sending heartbeats from the virtual machine to VCS in the physical host. You must increase the timeout value for the hearbeats.

Workaround

To increase the timeout value of the heartbeat:

1. Log into the physical host.

   Use the following command to modify the “IAATimeout” value for the virtual machine:

   ```
   # hasys -modify <GuestName> IAATimeout HigherTimeoutValue
   ```

   Where, GuestName refers to the hostname visible from within the virtual machine, and HigherTimeoutValue is a new considerably higher timeout value in seconds.

A virtual machine is unable to connect to VCS node

If a virtual machine is not connected to the physical host, and if heartbeats between the two are enabled, then the virtual machine tries to connect to the VCS node after every 30 seconds. If no VCS node exists, the following message appears:

```
No VCS server available for connection
```

Workaround

To ensure connectivity between the virtual machine and VCS, perform the following steps:

- Ensure that the VCS node (physical host) associated with the virtual machine is enabled to support ApplicationHA
- Ensure that the virtual machine is added to the ApplicationHA private VLAN. For more information:
  Viewing configuration details of a virtual machineSee "Viewing configuration details of a virtual machine" on page 54.

Disconnected virtual machine does not immediately appear faulted

When there is loss of connectivity between the virtual machine and physical host, VCS sets the connection status to NOT RESPONDING. The virtual machine status is moved to Faulted only after the ResponseTimeout period has lapsed.

Workaround

You can reset the ResponseTimeout attribute by executing the following steps:

1. Login to the physical host.
2. Run the following command:

   ```
   # hasys –modify GuestName ResponseTimeout  lower_timeout_value
   ```

   Where GuestName is the host name of the physical host and lower_timeout_value is the new, lower timeout value that you want to set.

Configured virtual machine appears to be disconnected from VCS

Even if you configure an application for monitoring on a virtual machine, and enable application heartbeat, the virtual machine appears to be not connected to the VCS cluster (physical host).

Workaround

1. Check the connection status after some time. The communication between ApplicationHA on the virtual machine and VCS on physical host may take some time to be detected by the VCS cluster.
2. Ensure that VCS on the physical host is enabled to support ApplicationHA. To view the configuration status of the associated physical host:

   See "Viewing configuration details of a virtual machine" on page 54.
3. On the physical host, perform the following verification steps:

- Verify that the clusext process `/opt/VRTSvcs/bin/clusext -core` is running. If this process is not running, wait for VCS to automatically start the process.
- If the clusext process is running, verify that the firewall on port 14142 is open.
- Verify that the VCS cluster level attribute EnableVMAutoDiscovery is set to 1.
- Verify that the vcsbr device is up with an IP address. If the device is not up, check the state of the resource "VCSVLANSG_IP". If the state is OFFLINE or FAULTED, bring the service group that contains the resource online.
- Verify that the DHCP server is started. If the DHCP server is not running, check the state of the VCSVLANSG_DHCP resource on the system. If the state is OFFLINE or FAULTED, bring online the service group that contains the resource.
- If the autoregistration mechanism is VirtIO, verify that SELinux is disabled on the physical host. Refer to Red Hat Linux documentation for information on how to disable SELinux.
- Verify that the libvirtd daemon is running on the VCS cluster node. You can use the `service libvirtd restart` command to start the daemon.

4. On the virtual machine, perform the following verification steps:
   - Verify that the VCSAppMonCERes resource inside the virtual machine is online. If this resource is not online, bring online the service group that contains this resource.
   - If the autoregistration mechanism for the virtual machines is VirtIO, verify that SELinux is disabled on the virtual machine. You must restart the virtual machine after disabling SELinux.
   - Verify that one of the interfaces on the virtual machine has an IP address plumbed from the network that was selected when VCS was configured to support ApplicationHA. If no such interface is visible on the virtual machine, other than the physical interface, it is possible that the interface was not hot plugged into virtual machine.
     You must restart the virtual machine from the physical host for the IP address to be plumbed.

Index