

Veritas High Availability 7.4.2 Solution Guide for VMware - Linux

Last updated: 2020-06-01

Legal Notice

Copyright © 2020 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054
<http://www.veritas.com>

Technical Support

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

<https://www.veritas.com/support>

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

<https://sort.veritas.com/documents>

Documentation feedback

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

infoscaledocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing the Veritas High Availability solution for VMware	7
	How the Veritas High Availability solution works in a VMware environment	7
	How the Veritas HA Plug-in for vSphere Web Client works with VCS	8
	How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host	10
	Getting started with vCenter-integrated Veritas High Availability solution	11
	Getting started with the VIOM-integrated Veritas High Availability solution	13
	Understanding Veritas High Availability terminology	14
	Supported VMware versions	16
	Important release information	16
Chapter 2	Deploying the Veritas High Availability solution	17
	Managing storage	17
	Installing the Veritas High Availability guest components	18
	Upgrading Veritas High Availability guest components	18
Chapter 3	Administering application availability from the vSphere Client	20
	Accessing the Veritas High Availability view	20
	Administering application monitoring from the Veritas High Availability view	21
	Understanding the Veritas High Availability view	22
	Configuring a cluster by using the VCS cluster configuration wizard	25
	To configure or unconfigure application monitoring	28
	Adding a system to a VCS cluster	29
	To start or stop applications	32
	To switch an application to another system	34

To add or remove a failover system	34
To suspend or resume application monitoring	39
To clear Fault state	40
To resolve a held-up operation	40
To determine application state	41
To remove all monitoring configurations	41
To remove VCS cluster configurations	41
Administering application monitoring settings	42
Administering application availability using Veritas High Availability	
dashboard	43
Understanding the dashboard work area	44
Accessing the dashboard	47
Monitoring applications across a data center	49
Monitoring applications across an ESX cluster	49
Searching for application instances by using filters	49
Selecting multiple applications for batch operations	50
Starting an application using the dashboard	50
Stopping an application by using the dashboard	50
Entering an application into maintenance mode	51
Bringing an application out of maintenance mode	52
Switching an application	52
Resolving dashboard alerts	53

Appendix A Roles and privileges 54

About the roles and privileges assigned in vCenter	54
Assigning customized privileges to VMwareDisks agent	55
About assigning privileges to VMwareDisks agent	55
Creating a role with customized privileges for VMwareDisks agent	
.....	56
Creating an ESX user account	56
Integrating an ESX user account with Active Directory	57
Assigning a role to an ESX user account	57

Appendix B Troubleshooting 59

Agent logging on virtual machine	59
Troubleshooting wizard-based configuration issues	60
Veritas High Availability Configuration wizard displays the	
"hadiscover is not recognized as an internal or external	
command" error	60
Running the 'hastop -all' command detaches virtual disks	60
Validation may fail when you add a failover system	60

Adding a failover system may fail if you configure a cluster with communication links over UDP	61
Troubleshooting issues with the Veritas High Availability view	61
Veritas high availability view is not visible from a cluster system	61
Veritas High Availability view does not display the application monitoring status	62
Veritas High Availability view may freeze due to special characters in application display name	62
If the Console host abruptly restarts, the high availability view may disappear	63
Veritas high availability view may fail to load or refresh	63
Operating system commands to unmount resource may fail	64

Introducing the Veritas High Availability solution for VMware

This chapter includes the following topics:

- [How the Veritas High Availability solution works in a VMware environment](#)

How the Veritas High Availability solution works in a VMware environment

The Veritas High Availability solution for VMware employs Cluster Server (VCS) and its agent framework to monitor the state of applications and their dependent components running on the virtual machines that use non-shared storage. Specific agents are available to monitor the application, storage, and network components. Together, these agents monitor the overall health of the configured applications by running specific commands, tests, or scripts.

The storage configuration in the VMware virtual environment determines how VCS functions differently in a non-shared virtual environment. The non-shared storage configuration in the VMware virtual environment involves the VMware VMDK and RDM disks that reside on the shared datastore. This datastore is accessible to multiple virtual machines. However, the disks are attached to a single virtual machine at any given point of time. VCS provides a new storage agent “VMwareDisks” that communicates with the VMware ESX/ESXi hosts to perform the disk detach and attach operations to move the storage disk between the virtual machines, in a VCS cluster.

Note: By default the VMwareDisks agent communicates with the ESX/ESXi host to perform the disk detach and attach operations. However, instead of the ESX/ESXi hosts you can choose to communicate with the vCenter Server to perform these operations.

See [“How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host”](#) on page 10.

In the event of an application failure, the agents attempt to restart the application services and components for a configurable number of times. If the application fails to start, they initiate an application failover to the failover target system. During the failover, the VMwareDisks agent moves the storage disk to the failover target system, the network agents bring the network components online, and the application-specific agents then start the application services on the failover target system.

In case of a virtual machine fault, the VCS agents begin to fail over the application to the failover target system. The VMwareDisks agent sends a disk detach request to the ESX/ESXi host. After the detach operation is successful, the agent proceeds to attach the disks to the new failover target system.

In a scenario where the ESX/ESXi host itself faults, the VCS agents begin to fail over the application to the failover target system that resides on another host. The VMwareDisks agent communicates with the new ESX/ESXi host and initiates a disk detach operation on the faulted virtual machine. The agent then attaches the disk to the new failover target virtual machine.

For details on the VCS configuration concepts and clustering topologies, refer to the *Cluster Server Administrator's Guide*.


For details on the application agents, refer to the application-specific agent guide.

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

How the Veritas HA Plug-in for vSphere Web Client works with VCS

The following graphic illustrates how the Veritas HA Plug-In for vSphere Web Client is deployed in a VMware virtual environment to support Cluster Server (VCS) tasks.

Figure 1-1 Deployment diagram



Note: The Veritas HA Plug-In for vSphere Web Client supports the high availability mode of the Veritas InfoScale Operations Manager Management Server (primary and secondary nodes). The plug-in also simultaneously supports multiple vCenter Servers. For simplicity, these components are not illustrated in the above diagram. Also, while the vSphere Web client support exists for multiple-node VCS clusters, for simplicity a single-node cluster of VCS is illustrated in this diagram.

The Veritas HA Plug-in for vSphere Web Client is distributed as an add-on for the Veritas InfoScale Operations Manager Management Server.

You must first install the add-on on the management server. You must also register the Veritas HA Plug-in for vSphere Web Client with all the required vCenter servers. As a result, the Veritas High Availability tab is inserted in the vSphere Web Client GUI.

If you add the VCS guests (virtual machines) to Veritas InfoScale Operations Manager Management Server as managed hosts, the management server discovers the health status of the applications monitored by VCS. The Management Server relays this information to the vCenter Server. This information is displayed in the Veritas High Availability tab of the vSphere Web Client.

Inside the Monitor tab of the vSphere Web Client, the Veritas High Availability tab is inserted. If you navigate to a virtual machine in the Web Client, and click Veritas High Availability tab, the Veritas High Availability View displays health information for the monitored application running on that virtual machine.

Similarly, if you navigate to an ESX cluster or datacenter, and click the tab, the Veritas High Availability Dashboard displays aggregate health information for all the monitored applications running in that cluster or datacenter.

You can perform application administration tasks from both the views. This includes starting and stopping the monitored applications, among other tasks. For details, see the *Veritas High Availability Solution Guide for VMware*.

How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host

In addition to the ESX hosts the VMwareDisks agent can also communicate the disk detach and attach operations with the vCenter Server to which the virtual machines belong.

In this scenario, in event of a failure, the VMwareDisks agent sends the disk detach and attach requests to the vCenter Server (instead of the ESX hosts). The vCenter Server then notifies the ESX host for these operations. Since the communication is directed through the vCenter Server, the agent successfully detaches and attaches the disks even if the ESX host and the virtual machines reside in a different network.

In a scenario where the host ESX/ESXi itself faults, the VMWareDisks agent from the target virtual machine sends a request to the vCenter Server to detach the disks from the failed virtual machine. However, since the host ESX has faulted, the request to detach the disks fails. The VMwareDisks agent from the target virtual machine now sends the disk attach request. The vCenter Server then processes this request and disks are attached to the target virtual machine. The application availability is thus not affected.

Limitation

The configuration of VMwareDisks agent to communicate with the vCenter Server has the following limitation:

If VMHA is not enabled and the host ESX faults, then even after the disks are attached to the target virtual machine they remain attached to the failed virtual machine. This issue occurs because the request to detach the disks fails since the host ESX itself has faulted. The agent then sends the disk attach request to the vCenter Server and attaches the disks to the target virtual machine.

Even though the application availability is not impacted, the subsequent power ON of the faulted virtual machine fails. This issue occurs because of the stale link between the virtual machine and the disks attached. Even though the disks are now attached to the target virtual machine the stale link with the failed virtual machine still exists.

Workaround

As a workaround, you must manually detach the disks from the failed virtual machine and then power ON the machine.

About the vCenter Server user account privileges

You must have the administrative privileges or must be a root user to communicate the disk detach and attach operations through the vCenter Server. If the vCenter Server user account fails to have the administrative privileges or is not a root user, then the disk detach and attach operation may fail, in event of a failure.

If you do not want to use the administrator user account or the root user, then you must create a role and add the following privileges to the created role:

- "Low level file operations" on datastore
- "Add existing disk" on virtual machine
- "Change resource" on virtual machine
- "Remove disk" on virtual machine

After you create a role and add the required privileges, you must add a local user to the created role. You can choose to add an existing user or create a new user.

Refer to the VMware product documentation for details on creating a role and adding a user to the created role.

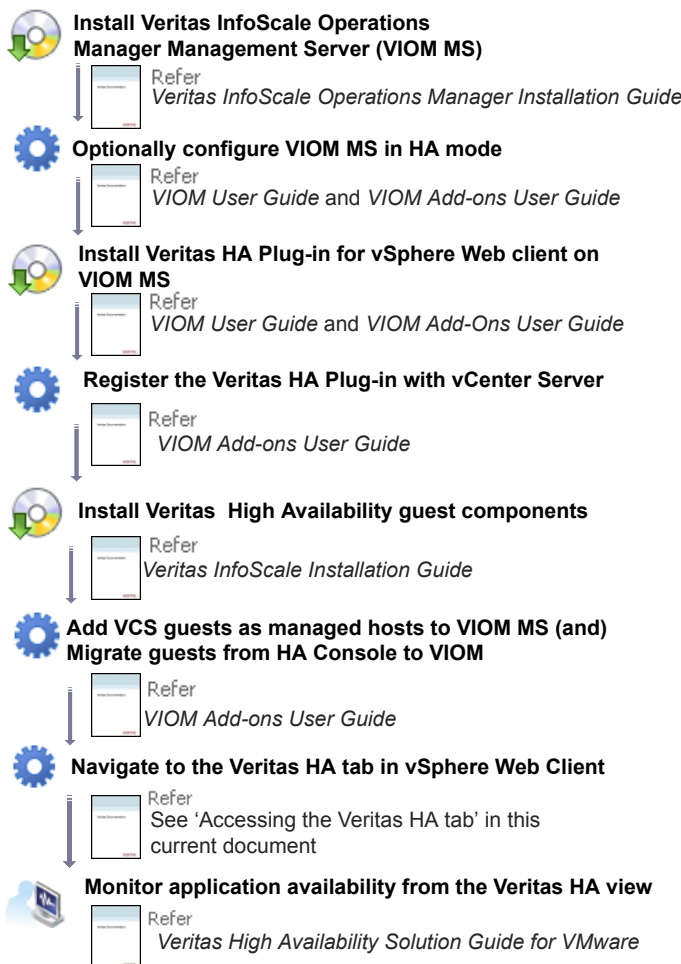
Getting started with vCenter-integrated Veritas High Availability solution

This section describes how to deploy the Veritas High Availability solution in a VMware environment. You can administer Veritas High Availability solutions over vSphere Web client.

Getting started with vSphere Web Client support

The vSphere Web Client support involves the installation and configuration of Veritas InfoScale Operations Manager Management Server and its Veritas HA Plug-in for VMware, apart from VCS and its agents.

The following figure represents the workflow for getting started with the Veritas High Availability solution and the corresponding document you must refer for details.



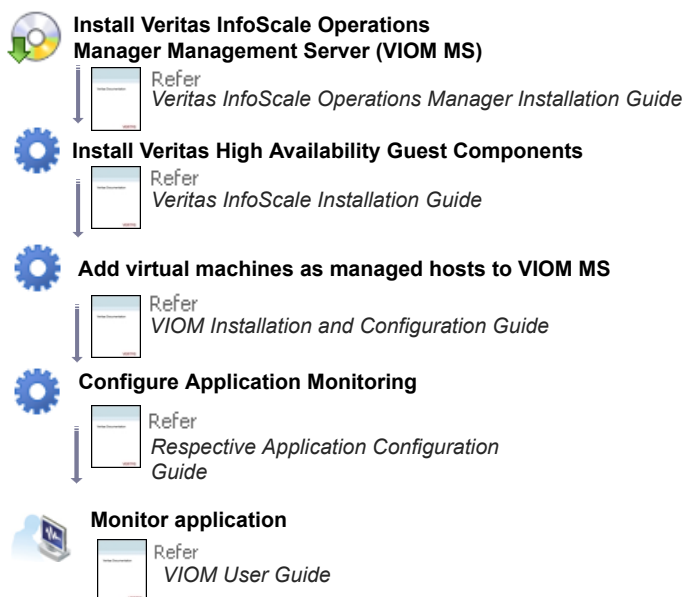
Note: If you use Veritas High Availability Console (version 6.2 or earlier) to administer VCS or ApplicationHA from vSphere Client (desktop), you need to upgrade to vSphere Web Client. You must migrate guest virtual machines from the Veritas High Availability Console to Veritas InfoScale Operations Manager. For details see tech note: https://www.veritas.com/support/en_US/article.TECH222796 and also refer *Veritas InfoScale Operations Manager Add-ons User Guide*. However, license management and vCenter integrated installation are not supported from the vSphere Web Client.

Getting started with the VIOM-integrated Veritas High Availability solution

The Veritas High Availability Solution operations can be integrated with Veritas InfoScale Operations Manager Management Server (VIOM MS) GUI by performing a few simple steps. The following figure lists the steps that you must perform to be able to administer the Veritas High Availability solution via VIOM MS GUI, and the documents that you may need to refer for further details.

Before you administer the solution over the VIOM MS GUI, also read the following section:

See [“About administering high availability with VOM”](#) on page 13.



About administering high availability with VOM

This document describes VMware vCenter-integrated steps to install, configure, and administer the Veritas High Availability solution. If you want to use the Veritas InfoScale Operations Manager (VOM)-integrated methods, you must be familiar with Storage Foundation and High Availability (SFHA) or Cluster Server (VCS) concepts and procedures. You must also be familiar with VOM concepts and procedures. You can download SFHA, VCS, and VOM documentation from the [SORT](#) website:

You can perform the following configuration and administration steps from the VOM Client:

Note: You cannot use the VOM client to install VCS or SFHA on guest virtual machines in a VMware environment. You can use the installer-based or other methods. For details, see the product-specific installation guide.

- Add virtual machines as managed hosts to VOM Management Server.
- Organize managed hosts
- Configure application monitoring
- Monitor configured applications
- Manage product licenses
- Manage user access (Role-based access control)
 For details on roles and privileges that you must assign to the user on VOM Management Server, see VOM documentation.

You cannot use the VOM client to administer the high availability of the virtualization infrastructure (ESX clusters) in a VMware environment. VMwareHA provides the high availability of ESX clusters.

Understanding Veritas High Availability terminology

The Veritas High Availability Solution for VMware refers to the integration of Cluster Server (VCS) tasks and functionalities with VMware vSphere GUI. As part of this integration, some terms have changed. [Table 1-1](#) lists such terms, along with their equivalent VCS term.

Based on context, the indicated synonyms have been used in the vSphere-integrated GUI, as well as in this document.

Table 1-1 Terminology disambiguation chart

Veritas High Availability Solutions documentation	Cluster Server documentation
Application	Service group
Application components	Resources
Veritas High Availability Guest Components	VCS, along with related RPMs, such as certain application-specific high availability agents
Application dependency	Service group dependency
Component dependency	Resource dependency

Table 1-1 Terminology disambiguation chart (*continued*)

Veritas High Availability Solutions documentation	Cluster Server documentation
System (until it becomes a part of a VCS cluster)	Node
OR	OR
VCS cluster system (once it joins a cluster)	VCS Cluster node
OR	OR
Failover target system (once it is included in the list of failover targets for an application)	Cluster member

Understanding operation names

Table 1-2 describes lists the names of VCS administrative operations/tasks that you can perform from the VMware vSphere Client GUI, as well as the equivalent operation names used in Cluster Server documentation:

Table 1-2 Task disambiguation chart

vSphere Client GUI-based operations	VCS operations
Start Application	Online Service Group
Stop Application	Offline Service Group
Switch	Switch To
Add Failover System	Add Node
Remove Failover System	Remove Node
Enter Maintenance Mode	Freeze Service Group
Exit Maintenance Mode	Unfreeze Service Group
Clear Fault State	Clear Fault
Resolve a Held-up Operation	Flush
Unconfigure Application Monitoring	Delete Service Group
Determine Application State	Probe
Stop/Start dependent components in order	Propagate (option for online/offline SGs)
Suspend application monitoring after reboot	Persistent (option for freeze)

Supported VMware versions

Refer the following support matrix for the VMware servers and management clients supported by Veritas High Availability solution 7.4.2:

https://www.veritas.com/support/en_US/article.000126342

Important release information

Review the Release notes for the latest information before you install the product.

Review the current compatibility lists to confirm the compatibility of your hardware and software:

- For important updates regarding this release, review the Late-Breaking News Tech Note on SORT:
https://www.veritas.com/support/en_US/article.000126340
- For the latest patches available for this release, go to:
<https://sort.veritas.com>
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware, visit the following URL:
https://www.veritas.com/support/en_US/article.000126344
- For the latest information on supported software, visit the following URL:
https://www.veritas.com/support/en_US/article.000126342

Deploying the Veritas High Availability solution

This chapter includes the following topics:

- [Managing storage](#)
- [Installing the Veritas High Availability guest components](#)
- [Upgrading Veritas High Availability guest components](#)

Managing storage

Configure the storage disks to save the application data.

VMware virtualization manages the application data by storing it on SAN LUNs (RDM file), or creating virtual disks on a local or networked storage attached to the ESX host using iSCSI, network, or Fibre Channel. The virtual disks reside on a datastore or a raw disk that exists on the storage disks used.

For more information, refer to the VMware documentation.

The application monitoring configuration in a VMware environment requires you to use the RDM or VMDK disk formats. During a failover, these disks can be deported from a system and imported to another system.

Consider the following to manage the storage disks:

- Use a networked storage and create virtual disks on the datastores that are accessible to all the ESX servers that hosts the VCS cluster systems.
- In case of virtual disks, create non-shared virtual disks (Thick Provision Lazy Zeroed).
- Add the virtual disks to the virtual machine on which you want to start the configured application.

- Create either LVM logical volumes or VxVM volumes.
- Mount the volumes on the mount point.

The following VCS storage agents are used to monitor the storage components involving non-shared storage:

- If the storage is managed using LVM, the LVMVolumeGroup and LVMLogicalVolume agents are used.
- If the storage is managed using VxVM, the DiskGroup and Volume agents are used.

Before configuring the storage, you can review the resource types and attribute definitions of these VCS storage agents. For details refer to the *Cluster Server Bundled Agents Reference Guide*.

Installing the Veritas High Availability guest components

To configure an application for high availability, you must install the Veritas High Availability guest components on the systems where you want to configure the application.

Veritas High Availability guest components, which include VCS component, application-specific high availability agents, and other related RPMs, are automatically installed as part of Veritas InfoScale installation. For details, see the *Veritas InfoScale Installation Guide*.

Note: Installation of Veritas High Availability guest components from vSphere Client menu is not supported in this release.

To administer application availability from the VMware vSphere Web Client, you must install Veritas InfoScale Operations Manager (VOM) Management Server. You must further install the Veritas HA Plug-in for vSphere Web Client on the VOM Management Server. For details, see Veritas InfoScale Operations Manager documentation.

Upgrading Veritas High Availability guest components

Upgrading Veritas High Availability guest components involves upgrading Cluster Server (VCS) and other allied components on the guests. If you have already

installed VCS on the guest virtual machines in a VMware environment, you can upgrade to release 7.2. For further details, see *VCS Configuration and Upgrade Guide*.

Administering application availability from the vSphere Client

This chapter includes the following topics:

- [Accessing the Veritas High Availability view](#)
- [Administering application monitoring from the Veritas High Availability view](#)
- [Administering application monitoring settings](#)
- [Administering application availability using Veritas High Availability dashboard](#)

Accessing the Veritas High Availability view

You can use one of the following methods to access the Veritas High Availability view.

- [To access the view from the vSphere Web Client](#)
- [To access the view from a browser window](#)

In the vSphere Web client, a prompt appears if single sign-on is not configured between Veritas Infoscale Operations Manager Management Server and the selected system.

If the system is not part of a VCS Cluster, the **Configure a VCS Cluster** link appears.

If you have not configured application monitoring, the **Configure an application for high availability** link appears.

To access the view from the vSphere Web Client

- 1 From the vSphere Web Client Home page, click the required vCenter Server.
- 2 Under Inventory trees, navigate to the required virtual machine.
- 3 In the Monitor tab, click **Veritas High Availability**.

The Veritas High Availability Health View for the selected virtual machine appears.

To access the view from a browser window

- ◆ Open a browser window and navigate to the following URL:

`https://SysNameorIP:5634/vcs/admin/application_health.html`

SysNameorIP is the system name or IP address of the system where you want to configure application monitoring.

Administering application monitoring from the Veritas High Availability view

The Veritas High Availability view is a Web-based graphic user interface that you can use to administer application availability with Cluster Server (VCS).

You can launch the Veritas High Availability view from an Internet browser, by using the following URL:

`https://VirtualMachineNameorIP:5634/vcs/admin/application_health.html?priv=ADMIN`

where *VirtualMachineNameorIP* is the system name or the IP address of the virtual machine from where you want to access the tab.

In a VMware virtual environment, you can embed the Veritas High Availability view as a tab in the vSphere Client menu (both the desktop and Web versions).

For details:

Note: You can administer application monitoring in two ways. One, using the Veritas High Availability tab as described below, and two, using the Veritas High Availability dashboard. Using the Veritas High Availability dashboard, you can administer application monitoring for multiple applications on multiple systems in a data center or ESX cluster. For more information about the dashboard-based operations:

See [“Administering application availability using Veritas High Availability dashboard”](#) on page 43.

Use the Veritas High Availability view to perform the following tasks:

- To configure a VCS cluster
- To add a system to a VCS cluster
- To configure and unconfigure application monitoring
- To unconfigure the VCS cluster
- To start and stop configured applications
- To add and remove failover systems
- To enter and exit maintenance mode
- To switch an application
- To determine the state of an application (components)
- To clear Fault state
- To resolve a held-up operation
- To modify application monitoring settings
- To view application dependency
- To view component dependency

Understanding the Veritas High Availability view

The Veritas High Availability view displays the consolidated health information for applications running in a Cluster Server (VCS) cluster. The cluster may include one or more systems.

The Veritas High Availability tab displays application information for the entire VCS cluster, not just the local system.


Note: If the local system is not part of any VCS cluster, then the Veritas Application High Availability view displays only the following link: **Configure a VCS cluster**.

If you are yet to configure an application for monitoring in a cluster of which the local system is a member, then the Veritas Application High Availability view displays only the following link: **Configure an application for high availability**.

The Veritas High Availability view uses icons, color coding, dependency graphs, and tool tips to report the detailed status of an application.

The Veritas High Availability view displays complex applications, in terms of multiple interdependent instances of that application. Such instances represent component groups (also known as service groups) of that application. Each service group in turn includes several critical components (resources) of the application.

The following figure shows the Veritas High Availability view, with one instance of Oracle Database and one instance of a generic application configured for high availability in a two-node VCS cluster:



1. Title bar
2. Actions menu
3. Aggregate Status Bar
4. Application dependency graph
5. Application table
6. Application-specific task menu
7. Component dependency graph

The Veritas High Availability view includes the following parts:

- Title bar: Displays the name of the VCS cluster, the Actions menu, the Refresh icon, the Alert icon. Note that the Alert icon appears only if the Veritas High Availability view fails to display a system, or displays stale data.
- Actions menu: Includes a drop-down list of operations that you can perform with effect across the cluster. These include: Configuring a cluster, Configuring an application for high availability; Unconfigure all applications; and Unconfigure VCS cluster.
- Aggregate status bar: Displays a summary of applications running in the cluster. This includes the total number of applications, and a breakdown of the number of applications in Online, Offline, Partial, and Faulted states.
- Application dependency graph: Illustrates the order in which the applications or application instances, must start or stop.

If an application must start first for another application to successfully start, the earlier application appears at a lower level in the graph. A line connects the two applications to indicate the dependency. If no such dependency exists, all applications appear in a single horizontal line.

- **Application table:** Displays a list of all applications configured in the VCS cluster that is associated with the local system.
Each application is listed in a separate row. Each row displays the systems where the application is configured for monitoring.
The title bar of each row displays the following entities to identify the application or application instance (service group):
 - Display name of the application (for example, Payroll application)
 - Type of application (for example, Custom)
 - Service group name
- **Application-specific task menu:** Appears in each application-specific row of the application table. The menu includes application-specific tasks such as Start, Stop, Switch, and a dropdown list of more tasks. The More dropdown list includes tasks such as Add a failover system, and Remove a failover system.
- **Component dependency graph:** Illustrates the order in which application components (resources) must start or stop for the related application or application instance to respectively start or stop. The component dependency graph by default does not appear in the application table. To view the component dependency graph for an application, you must click a system on which the application is running.
The track pad, at the right-bottom corner helps you navigate through component dependency graphs.
If you do not want to view the component dependency graph, in the top left corner of the application row, click **Close**.

To view the status of configured applications

In the application dependency graph, click the application for which you want to view the status. If the appropriate row is not already visible, the application table automatically scrolls to the appropriate row. The row displays the state of the application for each configured failover system in the cluster for that application.

If you click any system in the row, a component dependency graph appears. The graph uses symbols, color code, and tool tips to display the health of each application component. Roll the mouse over a system or component to see its health details.

The health of each application/application component on the selected system is displayed in terms of the following states:

Table 3-1 Application states

State	Description
Online	<p>Indicates that the configured application or application components are running on the virtual machine.</p> <p>If the application is offline on at least one other failover system, an alert appears next to the application name.</p>
Offline	<p>Indicates that the configured application or its components are not running on the virtual machine.</p>
Partial	<p>Indicates that either the application or its components are being started on the virtual machine or VCS was unable to start one or more of the configured components</p> <p>If the application is offline on at least one other failover system, an alert appears next to the application name.</p>
Faulted	<p>Indicates that the configured application or its components have unexpectedly stopped running.</p>

Configuring a cluster by using the VCS cluster configuration wizard

Perform the following steps to configure a Cluster Server (VCS) cluster by using the VCS Cluster Configuration wizard.

To configure a VCS cluster

- 1 Launch the VCS Cluster Configuration wizard.
- 2 Review the information on the Welcome panel and click **Next**.
The Configuration Inputs panel appears.
The local system is by default selected as a cluster system.
- 3 If you do not want to add more systems to the cluster, skip this step.
To add a system to the cluster, click **Add System**.

- 4 In the Add System dialog box, specify the following details for the system that you want to add to the VCS cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account you specified.
Use the specified user account on all systems	Select this check box to use the specified user account on all the cluster systems that have the same user name and password.

- 5 On the Configuration Inputs panel, do one of the following actions:
 - To add another system to the cluster, click Add System and repeat step 4.
 - To modify the specified User name or Password for a cluster system, use the edit icon.
 - Click **Next**
- 6 If you do not want to modify the security settings for the cluster, click **Next**, and proceed to step 8.

By default, the wizard configures single sign-on for secure cluster communication. If you want to modify the security settings for the cluster, click **Advanced Settings**.

- 7 In the Advanced settings dialog box, specify the following details and click **OK**.

Use Single Sign-on	Select to configure single sign-on using VCS Authentication Service for cluster communication. This option is enabled by default.
Use VCS user privileges	Select to configure a user with administrative privileges to the cluster. Specify the user name and password and click OK.

- 8
- On the Network Details panel, select the type of network protocol to configure the VCS cluster network links (Low Latency Transport or LLT module), and then specify the adapters for network communication.

The wizard configures the VCS cluster communication links using these adapters. You must select a minimum of two adapters per cluster system.

Note: By default, the LLT links are configured over Ethernet.

Select **Use MAC address for cluster communication (LLT over Ethernet)** or select **Use IP address for cluster communication (LLT over UDP)**, and specify the following details for each cluster system.

- To configure LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- To configure LLT over UDP, select the type of IP protocol (IPv4 or IPv6), and then specify the required details for each communication link.

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.
Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The second link is configured as a high-priority link.

To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: Veritas recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal VCS cluster communication over the link.

- 9 On the Configuration Summary panel, specify a cluster name and unique cluster ID and then click **Validate**.

Note: If multiple clusters exist in your network, the wizard validates if the specified cluster ID is a unique cluster ID among all clusters accessible from the current system. Among clusters that are not accessible from the current system, you must ensure that the cluster ID you specified is unique

- 10 Review the VCS Cluster Configuration Details and then click **Next** to proceed with the configuration
- 11 On the Implementation panel, the wizard creates the VCS cluster.
 The wizard displays the status of the configuration task. After the configuration is complete, click **Next**.
 If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to configure the VCS cluster.
- 12 On the Finish panel, click **Finish** to complete the wizard workflow.
 This completes the VCS cluster configuration.

To configure or unconfigure application monitoring

Use the Veritas High Availability view to configure or un-configure an application for monitoring in a cluster under Cluster Server (VCS) control.

The view provides you with specific links to perform the following configuration tasks:

- Configure a VCS cluster:
 If the local system is not part of a VCS cluster, the Veritas High Availability view appears blank, except for the link **Configure a VCS cluster**.
 Before you can configure application monitoring, you must click the **Configure a VCS cluster** link to launch the VCS Cluster Configuration wizard. The wizard helps you configure a VCS cluster, where the local system is by default a cluster system.
- Configure the first application for monitoring in a VCS cluster:

If you have not configured any application for monitoring in the cluster, the Veritas High Availability view appears blank except for the link **Configure an application for high availability**.

Click the link to launch the Veritas High Availability Configuration Wizard. Use the wizard to configure application monitoring.

Also, in applications where the Veritas High Availability Configuration Wizard is supported, for detailed wizard-based steps, see the application-specific VCS agent configuration guide. For custom applications, see the *Cluster Server Generic Application Agent Configuration Guide*.

- Add a system to the VCS cluster:
Click **Actions > Add a System to VCS Cluster** to add a system to a VCS cluster where the local system is a cluster member. Adding a system to the cluster does not automatically add the system as a failover system for any configured application. To add a system as a failover system, see (add a cross-reference to 'To add or remove a failover system'.
- Configure an application or add an application to the existing application monitoring configuration:
Click **Actions > Configure an application for high availability** to launch the Veritas High Availability Application Monitoring Configuration Wizard. Use the wizard to configure application monitoring.
- Unconfigure monitoring of an application:
In the appropriate row of the application table, click **More > Unconfigure Application Monitoring** to delete the application monitoring configuration from the VCS.
Note that this step does not remove VCS from the system or the cluster, this step only removes the monitoring configuration for that application.
Also, to unconfigure monitoring for an application, you can perform one of the following procedures: unconfigure monitoring of all applications, or unconfigure VCS cluster.
- Unconfigure monitoring of all applications:
Click **Actions > Unconfigure all applications**. This step deletes the monitoring configuration for all applications configured in the cluster.
- Unconfigure VCS cluster:
Click **Actions > Unconfigure VCS cluster**. This step stops the VCS cluster, removes VCS cluster configuration, and unconfigures application monitoring.

Adding a system to a VCS cluster

Perform the following steps to add a system to a Cluster Server (VCS) cluster by using the VCS Cluster Configuration wizard.

The system from where you launch the wizard must be part of the cluster to which you want to add a new system.

To add a system to a VCS cluster

- 1
- Access the Veritas High Availability view (for any system belonging the required cluster).
- 2
- Click **Actions > Add System to VCS Cluster**.
The VCS Cluster Configuration Wizard is launched.
- 3
- Review the information on the Welcome panel and click **Next**.
The Configuration Inputs panel appears, along with the cluster name, and a table of existing cluster systems.
- 4
- To add a system to the cluster, click **Add System**.
- 5
- In the Add System dialog box, specify the following details for the system that you want to add to the VCS cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account you specified.
Use the specified user account on all systems	Select this check box to use the specified user account on all the cluster systems that have the same user name and password.

- 6
- On the Configuration Inputs panel, do one of the following actions:
 - To add another system to the cluster, click **Add System** and repeat step 5.
 - To modify the User name or Password for a cluster system, use the edit icon.
 - Click **Next**
- 7
- On the Network Details panel, specify the adapters for network communication (Low Latency Transport or LLT module of VCS) for the system. The wizard configures the VCS cluster communication links using these adapters. You must select a minimum of two adapters.

Note: You cannot modify the existing type of cluster communication (LLT over Ethernet or LLT over UDP).

- If the existing cluster uses LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- If the existing cluster uses LLT over UDP, select the type of IP protocol (IPv4 or IPv6), and then specify the required details for each communication link.

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.
Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The other link is configured as a high-priority link.

To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: Veritas recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal VCS cluster communication over the link.

- 8
- On the Configuration Summary panel, review the VCS Cluster Configuration Details.

- 9 On the Implementation panel, the wizard creates the VCS cluster.

The wizard displays the status of the configuration task. After the configuration is complete, click **Next**.

If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to add the required system to the VCS cluster.

- 10 On the Finish panel, click **Finish** to complete the wizard workflow.

To start or stop applications

Use the following options on the Veritas High Availability view to control the status of the configured application and the associated components or component groups (application instances).

Note that the **Start** and **Stop** links are dimmed in the following cases:

- If you have not configured any associated components or component groups (resources or service groups) for monitoring
- If the application is in maintenance mode
- If no system exists in the cluster, where the application is not already started or stopped as required.

To start an application

- 1 In the appropriate row of the application table, click **Start**.
- 2 If the application (service group) is of the failover type, on the Start Application panel, click **Any system**. VCS uses pre-defined policies to decide the system where to start the application.

If the application (service group) is of the parallel type, on the Start Application panel, click **All systems**. VCS starts the application on all required systems, where the service group is configured.

Note: Based on service group type, either the Any system or the All Systems link automatically appears.

To learn more about policies, and parallel and failover service groups, see the *Cluster Server Administrator's Guide*.

If you want to specify the system where you want to start the application, click **User selected system**, and then click the appropriate system.

- 3 If the application that you want to start requires other applications or component groups (service groups) to start in a specific order, then check the **Start the dependent components in order** check box, and then click **OK**.

To stop an application

- 1 In the appropriate row of the application table, click **Stop**.
- 2 If the application (service group) is of the failover type, in the Stop Application Panel, click **Any system**. VCS selects the appropriate system to stop the application.

If the application (service group) is of the parallel type, in the Stop Application Panel click **All systems**. VCS stops the application on all configured systems.

Note: Based on service group type, either the Any system or the All Systems link automatically appears.

To learn more about parallel and failover service groups, see the *Cluster Server Administrator's Guide*.

If you want to specify the system where you want to stop the application, click **User selected system**, and then click the appropriate system.

- 3 If the application that you want to stop requires other applications or component groups (service groups) to stop in a specific order, then check the **Stop the dependent components in order** check box, and then click **OK**.

To switch an application to another system

If you want to gracefully stop an application on one system and start it on another system in the same cluster, you must use the Switch link. You can switch the application only to a system where it is not running.

Note that the Switch link is dimmed in the following cases:

- If you have not configured any application components for monitoring
- If you have not specified any failover system for the selected application
- If the application is in maintenance mode
- If no system exists in the cluster, where the application can be switched
- If the application is not in online/partial state on even a single system in the cluster

To switch an application

- 1 In the appropriate row of the application table, click **Switch**.
- 2 If you want VCS to decide to which system the application must switch, based on policies, then in the Switch Application panel, click **Any system**, and then click **OK**.

To learn more about policies, see the *Cluster Server Administrator's Guide*.

If you want to specify the system where you want to switch the application, click **User selected system**, and then click the appropriate system, and then click **OK**.

VCS stops the application on the system where the application is running, and starts it on the system you specified.

To add or remove a failover system

Each row in the application table displays the status of an application on systems that are part of a VCS cluster. The displayed system/s either form a single-system Cluster Server (VCS) cluster with application restart configured as a high-availability measure, or a multi-system VCS cluster with application failover configured. In the displayed cluster, you can add a new system as a failover system for the configured application.

The system must fulfill the following conditions:

- The system is not part of any other VCS cluster.
- The system has at least two network adapters.

- The host name of the system must be resolvable through the DNS server or, locally, using `/etc/hosts` file entries.
- The required ports are not blocked by a firewall.
- The application is installed identically on all the systems, including the proposed new system.

To add a failover system, perform the following steps:

Note: The following procedure describes generic steps to add a failover system. The wizard automatically populates values for initially configured systems in some fields. These values are not editable.

To add a failover system

- 1** In the appropriate row of the application table, click **More > Add Failover System**.
- 2** Review the instructions on the welcome page of the Veritas High Availability Configuration Wizard, and click **Next**.

- 3 If you want to add a system from the Cluster systems list to the Application failover targets list, on the Configuration Inputs panel, select the system in the Cluster systems list. Use the Edit icon to specify an administrative user account on the system. You can then move the required system from the Cluster system list to the Application failover targets list. Use the up and down arrow keys to set the order of systems in which VCS agent must failover applications.

If you want to specify a failover system that is not an existing cluster node, on the Configuration Inputs panel, click **Add System**, and in the Add System dialog box, specify the following details:

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user name with administrative privileges on the system. If you want to specify the same user account on all systems that you want to add, check the Use the specified user account on all systems box.
Password	Specify the password for the account you specified.
Use the specified user account on all systems	Click this check box to use the specified user credentials on all the cluster systems.

The wizard validates the details, and the system then appears in the Application failover target list.

- 4 If you are adding a failover system from the existing VCS cluster, the Network Details panel does not appear.

If you are adding a new failover system to the existing cluster, on the Network Details panel, review the networking parameters used by existing failover systems. Appropriately modify the following parameters for the new failover system.

Note: The wizard automatically populates the networking protocol (UDP or Ethernet) used by the existing failover systems for Low Latency Transport communication. You cannot modify these settings.

- To configure links over ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- To configure links over UDP, specify the required details for each communication link.

Network Adapter	<p>Select a network adapter for the communication links.</p> <p>You must select a different network adapter for each communication link.</p> <p>Veritas recommends that one of the network adapters must be a public adapter and the VCS cluster communication link using this adapter is assigned a low priority.</p> <p>Note: Do not select the teamed network adapter or the independently listed adapters that are a part of teamed NIC.</p>
IP Address	Select the IP address to be used for cluster communication over the specified UDP port.
Port	<p>Specify a unique port number for each link. You can use ports in the range 49152 to 65535.</p> <p>The specified port for a link is used for all the cluster systems on that link.</p>
Subnet mask	Displays the subnet mask to which the specified IP belongs.

- 5** If a virtual IP is not configured as part of your application monitoring configuration, the Virtual Network Details page is not displayed. Else, on the Virtual Network Details panel, review the following networking parameters that the failover system must use, and specify the NIC:

Virtual IP address	Specifies a unique virtual IP address.
Subnet mask	Specifies the subnet mask to which the IP address belongs.
NIC	For each newly added system, specify the network adaptor that must host the specified virtual IP.

- 6 If the newly added failover system is associated with a different ESX host as compared to other systems, then on Target ESX Details page, specify the ESX host of the newly added failover system. Also specify the administrative user account details associated with the ESX host.

Note: If the application for which you are adding a failover system does not use storage attached directly to the ESX host, the wizard does not display this page.

If the new failover system runs on a different ESX host, or is configured to failover to another ESX host, specify that ESX host. To specify the ESX host, click **Add ESX Host** and on the Add ESX Host dialogue box, specify the following details, and then click **Next**:

ESX hostname or IP address	Specify the target ESX hostname or IP address. The virtual machines can fail over to this ESX host during vMotion. Specify an ESX host that has the same mount points as those currently used by the application.
User name	Specify a user account for the ESX host. The user account must have administrator privileges on the specified ESX host.
Password	Specify the password associated with the user name you specified.

The wizard validates the user account and the storage details on the specified ESX host, and uses this account to move data disks during vMotion.

- 7 On the Configuration Summary panel, review the VCS cluster configuration summary, and then click **Next** to proceed with the configuration.
- 8 On the Implementation panel, the wizard adds the specified system to the VCS cluster, if it is not already a part. It then adds the system to the list of failover targets. The wizard displays a progress report of each task.
 - If the wizard displays an error, click **View Logs** to review the error description, troubleshoot the error, and re-run the wizard from the Veritas High Availability view.
 - Click **Next**.
- 9 On the Finish panel, click **Finish**. This completes the procedure for adding a failover system. You can view the system in the appropriate row of the application table.

Similarly you can also remove a system from the list of application failover targets.

Note: You cannot remove a failover system if an application is online or partially online on the system.

To remove a failover system

- 1 In the appropriate row of the application table, click **More > Remove Failover System**.
- 2 On the Remove Failover System panel, click the system that you want to remove from the monitoring configuration, and then click **OK**.

Note: This procedure only removes the system from the list of failover target systems, not from the VCS cluster. To remove a system from the cluster, use VCS commands. For details, see the *VCS Administrator's Guide*.

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, Cluster Server (VCS) 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, VCS freezes the application configuration.

The **Enter Maintenance Mode** link is automatically dimmed if the application is already in maintenance mode. Conversely, if the application is not in maintenance mode, the **Exit Maintenance Mode** link is dimmed.

The Veritas High Availability tab provides the following options:

To enter maintenance mode

- 1 In the appropriate row, click **More> Enter Maintenance Mode**.
 During the time the monitoring is suspended, Veritas high availability solutions do not monitor the state of the application and its dependent components. The Veritas High Availability view does not display the current status of the application. If there is any failure in the application or its components, VCS takes no action.
- 2 While in maintenance mode, if a virtual machine restarts, if you want application monitoring to remain in maintenance mode, then in the Enter Maintenance Mode panel, check the **Suspend the application availability even after reboot** check box, and then click **OK** to enter maintenance mode.

To exit the maintenance mode

- 1 In the appropriate row, click **More> Exit Maintenance Mode**, and then click **OK** to exit maintenance mode.
- 2 Click the Refresh icon in the top right corner of the Veritas High Availability view, to confirm that the application is no longer in maintenance mode.

To clear Fault state

When you fix an application fault on a system, you must further clear the application Faulted state on that system. Unless you clear the Faulted state, VCS cannot failover the application on that system.

You can use the Veritas High Availability view to clear this faulted state at the level of a configured application component (resource).

The Clear Fault link is automatically dimmed if there is no faulted system in the cluster.

To clear Fault state

- 1 In the appropriate row of the application table, click **More > Clear Fault state**.
- 2 In the Clear Fault State panel, click the system where you want to clear the Faulted status of a component, and then click **OK**.

To resolve a held-up operation

When you try to start or stop an application, in some cases, the start or stop operation may get held-up mid course. This may be due to VCS detecting an incorrect internal state of an application component. You can resolve this issue by using the resolve a held-up operation link. When you click the link, VCS appropriately resets the internal state of any held-up application component. This process prepares

the ground for you to retry the original start or stop operation, or initiate another operation.

To resolve a held-up operation

- 1 In the appropriate row of the application table, click **More > Resolve a held-up operation**.
- 2 In the Resolve a held-up operation panel, click the system where you want to resolve the held-up operation, and then click **OK**.

To determine application state

The Veritas High Availability view displays the consolidated health information of all applications configured for monitoring in a VCS cluster. The application health information is automatically refreshed every 60 seconds.

If you do not want to wait for the automatic refresh, you can instantaneously determine the state of an application by performing the following steps:

To determine application state

- 1 In the appropriate row of the Application table, click **More > Determine Application State**.
- 2 In the Determine Application State panel, select a system and then click **OK**.

Note: You can also select multiple systems, and then click **OK**.

To remove all monitoring configurations

To discontinue all existing application monitoring in a VCS cluster, perform the following step:

- On the Veritas High Availability view, in the Title bar, click **Actions > Unconfigure all applications**. When a confirmation message appears, click **OK**.

To remove VCS cluster configurations

If you want to create a different VCS cluster, say with new systems, a different LLT protocol, or secure communication mode, you may want to remove existing VCS cluster configurations. To remove VCS cluster configurations, perform the following steps:

Note: The following steps deletes all cluster configurations, (including networking and storage configurations), as well as application-monitoring configurations.

- On the Title bar of the Veritas High Availability view, click **Actions >Unconfigure VCS cluster**.
- In the Unconfigure VCS Cluster panel, review the Cluster Name and Cluster ID, and specify the User name and Password of the Cluster administrator, and then click **OK**.

Administering application monitoring settings

The Veritas High Availability view lets you define and modify settings that control application monitoring with Cluster Server(VCS). You can define the settings on a per application basis. The settings apply to all systems in a VCS cluster, where that particular application is configured for monitoring.

The following settings are available:

- **App.StartStopTimeout:** When you click the **Start Application** or **Stop Application**, or **Switch Application** links in the Veritas High Availability view, VCS initiates an application start or stop, respectively. This option defines the number of seconds that VCS must wait for the application to start or stop, after initiating the operation. You can set a value between 0 and 300 seconds for this attribute; the default value is 30 seconds.
If the application does not respond in the stipulated time, the tab displays an alert. The alert states that the operation may take some more time to complete and that you must check the status after some time. 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. VCS continues to wait for the application response even after the timeout interval elapses.
If the application fails to start or stop, VCS takes the necessary action depending on the other configured remedial actions.
- **App.RestartAttempts:** This setting defines the number of times that VCS must try to restart a failed application. The value of App.RestartAttempts may vary between 0 and 5; the default value is 0. If an application fails to start within the specified number of attempts, VCS fails over the application to a configured failover system.
- **App.DisplayName:** This setting lets you specify an easy-to-use display name for a configured application. For example, Payroll Application. VCS may internally use a different application name to uniquely identify the application. However, the internal string, for example OraSG2, may not be intuitive to understand, or

easy to recognize while navigating the application table in the Veritas High Availability view.

Moreover, once configured, you cannot edit the application name, while you can modify the application display name as required. Note that the Veritas High Availability view displays both the application display name and the application name.

Administering application availability using Veritas High Availability dashboard

The Veritas High Availability Dashboard is a consolidated graphic user interface that lets you administer application monitoring on systems in a VMware vCenter-administered data center.

The dashboard is fully integrated with the VMware vSphere Client GUI. The dashboard appears in the Veritas High Availability tab of the VMware vSphere Client GUI. To view the dashboard, select a data center or an ESX cluster in the inventory, and then click the Veritas High Availability tab.

Note: To administer application availability using the dashboard, single sign-on between the system and Veritas High Availability Console must be configured. Also, the application-specific agent must be appropriately configured.

For more information, see the *Veritas High Availability Solution Guide for VMware*.

On the dashboard, you can view the aggregate health statistics for monitored applications across a data center. You can also drill down to an ESX cluster and view monitored applications running in that cluster.

To understand how to navigate across the dashboard:

See [“Understanding the dashboard work area”](#) on page 44.

You can drill down to an individual application and perform the following administrative actions:


- Start application
- Stop application
- Enter maintenance mode
- Exit maintenance mode
- Switch application (to another system)

Understanding the dashboard work area

The Veritas High Availability dashboard displays the aggregate application health status information for a datacenter or an ESX cluster.

Depending on whether you click a datacenter or a VMware cluster in the inventory view (left pane) of the VMware vSphere Client GUI, the dashboard displays the aggregate application status information. Apart from the application table described in detail below, the dashboard uses color code and tool tips to indicate the status of an application.

The following figure illustrates the dashboard work area. Note that the red boxes highlight the key GUI elements:



In the above figure, the labels stand for the following elements of the dashboard

- | | | | | | |
|---|----------------------|---|------------------------|---|--------------------------|
| 1 | Aggregate status bar | 2 | ESX cluster/host table | 3 | Taskbar |
| 4 | Filters menu | 5 | Application table | 6 | Systems table (dropdown) |

Aggregate status bar

The aggregate status bar of the dashboard displays the following details:

- Number of ESX clusters that have applications configured for monitoring with VCS
- Number of configured applications in the selected data center

- Number of faulted applications
- Number of applications in partial state
- Number of online applications
- Number of offline applications

ESX cluster/host table

The Veritas High Availability dashboard displays this table only if you click a datacenter in the inventory view of the vSphere Client, and then click the Veritas High Availability tab.

The cluster table lists the following statistics per ESX cluster (or independent ESX host) in the data center:

- Number of configured applications
- Number of faulted applications
- Number of applications in partial state
- Number of online applications
- Number of offline applications
- Overall status (percentage of healthy applications)

If you click a row in the ESX cluster/host table, the application table of the dashboard displays monitored applications running on systems hosted by the selected ESX cluster or ESX host (an ESX server that is not part of an ESX cluster).

Note: This is the only method to navigate to applications running on systems hosted by standalone ESX hosts, by using the Veritas High Availability dashboard.

Taskbar

The taskbar displays icons for various administrative tasks. A tool tip highlights the task that each icon represents.

The dashboard supports the following tasks:

- Start Application: Starts a configured application
- Stop Application: Stops a configured application
- Enter Maintenance Mode: Suspends application monitoring of the configured application. In maintenance mode, VCS does not monitor the state of the application, and its dependent components.

- Exit Maintenance Mode: Resumes application monitoring for a configured application.
- Switch Application: Switches and an application gracefully from one system to another.

Filters menu

The filters menu lets you dynamically filter the applications that are displayed in the applications table. You can filter the applications by the following parameters:

- Application name
- Application status
- Search (by a string)

Application table

If you click an ESX cluster in the ESX cluster/host table, or in the inventory view of the VMware vSphere Client, then the list of applications running in that ESX cluster appears in the application table of the dashboard.

If you click an ESX host (an ESX server that is not part of an ESX cluster) in the ESX cluster/host table, then the list of applications that are configured on systems hosted by that ESX server appears. Note that this is the only route to navigate to such applications through the dashboard.

The following table lists each column in the application table and its description:

Column	Description
Applications	Indicates the application name.
Service Groups	<p>Indicates the group of critical application components that VCS uses to determine the health of a monitored application. Service group is a VCS term.</p> <p>VCS may use more than one service group to monitor a complex application. The dashboard displays each service group of such an application as a separate instance of that application.</p>

Column	Description
Status	<p>This column indicates the effective status of an application in a VCS cluster. It does not indicate the state of the application on per member system. For example, in a two-system cluster, if the application has faulted on one system but has failed over to another system, then this column states the state of the application as Online.</p> <p>Indicates one of the following states of an application:</p> <ul style="list-style-type: none"> ■ Online ■ Offline ■ Faulted ■ Partial <p>Note: After you perform an administrative task such as starting or stopping an application, or entering or exiting maintenance mode, it takes a few seconds for the dashboard to reflect the revised status of the configured application.</p>
Systems	<p>Indicates the number of systems where the application is configured for monitoring. To view more information about all such systems, click the (...) icon. The System table (dropdown) appears, listing the ESX host name of each configured system, the VM name (system name), and the status of the application on each system.</p>
Alerts and description	<p>Displays a triangular alert icon (!) and describes the reason for the alert. This column displays alerts in two cases: a) If the application status record is stale; b) If the application has faulted on a system.</p> <p>For stale records, the column includes the timestamp of the last received health record. In case of application fault, the column provides details of the system where the fault occurred.</p>

Accessing the dashboard

You can use the Veritas High Availability dashboard to perform one of the following actions:

- Identify all instances and failover systems of one or more applications running in a data center
- Drill down to a specific application, and perform an administrative action on the application
- View alerts for faulted applications and stale application health reports

Prerequisites for accessing the dashboard

Before you access the Veritas High Availability dashboard to administer an application, ensure:

- Single sign-on is configured between the Veritas High Availability Console and the systems hosting the monitored applications
- Veritas High Availability Console is able to communicate with Veritas High Availability guest components on designated port (port 5634).
- The application that you want to administer is configured for application monitoring with Veritas High Availability

How to access the dashboard

When you install Veritas High Availability guest components, the product installation script or wizard automatically installs the required dashboard components. As a result, the Veritas High Availability Dashboard appears in the **Veritas High Availability** tab of the vSphere Client.

You must, however, ensure that Veritas High Availability is successfully installed and that you have adequate user privileges to access the dashboard.

To access dashboard

Perform the following step:

- In the inventory view (left pane) of the vSphere Client, click a datacenter or a VMware cluster. In the right pane, to view the Veritas High Availability dashboard, click the Veritas High Availability tab.

Who can access the dashboard

To access High Availability dashboard, the VMware vCenter administrator must assign one the following roles to you:

- Guest: View application state
- Operator: View application state and perform administrative actions
- Admin: View application state and perform administrative actions. You can also configure application availability monitoring in this role, but not from the dashboard.

For more information on roles and privileges:

See [“About the roles and privileges assigned in vCenter”](#) on page 54.

Monitoring applications across a data center

If you click a data center in the inventory view of the VMware vSphere Client, and then click the Veritas High Availability tab, the dashboard appears, displaying the aggregate health information of applications running inside various ESX clusters.

You can use filters to drill down from all applications running across the data center and view a single application and its various instances in the data center.

Monitoring applications across an ESX cluster

If you click an ESX cluster in the inventory view of the VMware vSphere Client, and then click the tab, the dashboard displays the consolidated information on the systems and applications running in the ESX cluster. The dashboard also displays the application health and application monitoring information.

You can use filters to drill down from all applications running in the ESX cluster, to view a single application and its various instances in the ESX cluster.

Searching for application instances by using filters

The High Availability dashboard lets you search for all instances of a particular application in the selected datacenter or an ESX cluster. Various filters enable to search for the application that you want to monitor. You can use multiple filters simultaneously to search for an application.

The following table lists each field in the filter menu and its description:

Field	Description
Application	Lets you specify the name of the application that you want to filter in the application table. A drop-down list displays all the applications that are configured in the datacenter or ESX cluster. Click to select the name of the application that you want to filter.
Status	Lets you specify the status of the application by which you want to filter the application table. A drop-down list displays the following status values: Online, Offline, Faulted, and Partial.
Search	<p>Lets you search for an application by using a string or pattern of characters. Enter the string using which you want to filter applications. As you enter the string in the Search box, the dashboard dynamically filters the applications.</p> <p>Note: The dashboard searches for the specified string in the Systems column.</p>

Selecting multiple applications for batch operations

You can select one or more instances of an application for administering by using the dashboard as follows:

- To select one application instance, click inside the row of that application instance.
- To select various instances, keep the **Control** key pressed and then click inside the row of each instance.
- To select a batch of consecutive entries in the application table, keep the **Shift** key pressed, click inside the row of the first instance, and then click inside the row of the last instance. Alternatively, you can keep the **Shift** key pressed and drag the mouse to mark a block of consecutive entries.
- To select all instances in the application table, click **Select All**.

Starting an application using the dashboard

To start an application, perform the following steps in the application table of the dashboard.

To start an application

- 1 Filter the applications that you want to start.
 See [“Searching for application instances by using filters”](#) on page 49.
 The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to perform a batch operation.
 See [“Selecting multiple applications for batch operations”](#) on page 50.
- 3 To start the application, in the taskbar, click the appropriate icon (use the tool tip to recognize the appropriate icon).
- 4 In the Start Application panel, click the systems where you want to start the application. Note that you can start the application on any of the systems displayed for each application.
 Click **OK**.

Stopping an application by using the dashboard

To stop an application on one or more virtual machines, perform the following steps in the application table of the High Availability dashboard.

To stop an application

- 1 Filter the applications that you want to stop.
See [“Searching for application instances by using filters”](#) on page 49.
The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to perform a batch operation.
See [“Selecting multiple applications for batch operations”](#) on page 50.
- 3 To stop the application, in the taskbar, click the appropriate icon (use the tool tip to recognize the appropriate icon).
- 4 In the Stop Application panel, from the dropdown list, click the systems where you want to stop the application.
Click **OK**.

Entering an application into maintenance mode

You may need to intentionally take an application offline for maintenance purposes, without triggering a corrective response from Cluster Server (VCS).

To enter an application into maintenance mode, perform the following steps in the application table of the High Availability dashboard.

Note: The maintenance mode configuration is application-specific, not system-specific.

To enter maintenance mode

- 1 Filter the application that you want to gracefully take offline for maintenance.
See [“Searching for application instances by using filters”](#) on page 49.
The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to perform a batch operation.
See [“Selecting multiple applications for batch operations”](#) on page 50.
- 3 To enter maintenance mode, in the taskbar, click the appropriate icon for entering maintenance mode (use the tool tip to recognize the appropriate icon).

- 4 If a system restarts while the application is in maintenance mode, and you want the application to remain in maintenance mode, then in the Enter Maintenance Mode panel, check the **Suspend the application availability even after reboot**.
- 5 On the Enter Maintenance Mode panel, click **OK**.

Bringing an application out of maintenance mode

To bring an application out of maintenance mode on one or more systems, perform the following steps in the application table of the High Availability dashboard.

To exit maintenance mode

- 1 Filter the applications that you want to bring out of maintenance mode.
 See [“Searching for application instances by using filters”](#) on page 49.
 The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to bring out of maintenance mode.
 See [“Selecting multiple applications for batch operations”](#) on page 50.
- 3 To bring the applications out of maintenance mode, in the taskbar, click the appropriate icon for exiting maintenance mode (use the tool tip to recognize the appropriate icon).
- 4 In the Exit Maintenance Mode panel, click **OK**.

Switching an application

To gracefully switch an application from one system to another, perform the following steps in the application table of the dashboard.

Note: You can switch an application only if the application monitoring configuration includes one or more failover systems.

To switch an application

- 1 Filter the applications that you want to switch to another node.
 See [“Searching for application instances by using filters”](#) on page 49.
 The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to perform a batch operation.
 See [“Selecting multiple applications for batch operations”](#) on page 50.

- 3 To switch the applications, in the taskbar, click the appropriate icon (use the tool tip to recognize the appropriate icon).
- 4 In the Switch Application panel, select the systems where you want to switch the applications, and then click **OK**. Cluster Server takes the applications offline on the existing systems, and brings them online on the systems that you specified.

Resolving dashboard alerts

The Alerts and Description column in the application table of the High Availability dashboard marks application alerts with the alert (!) icon. This occurs in the following cases:

- Stale entries: Stale entries occur either due to a system (virtual machine) issues or connectivity issues. When this occurs, the system fails to send application heartbeats to the dashboard. If the system fails to send the heartbeat for two consecutive heartbeat intervals, the dashboard displays the alert icon.

Note: You can filter stale entries using the **Search** option and searching with the string "stale".

- Application faults: Application faults may occur due to reasons beyond Cluster Server (VCS) control, such as storage failure. In such cases, you must investigate and appropriately resolve the issue, and then clear the Faulted status of the application. To view only application fault alerts, in the Alerts and Description column, click the **Hide Old Entries** check box.

Note: It is important that you fix application faults, and then clear the Fault status. Else, the VCS cannot fail over applications to the faulted system, and application availability may be compromised. For more information, See ["To clear Fault state"](#) on page 40.

Roles and privileges

This appendix includes the following topics:

- [About the roles and privileges assigned in vCenter](#)
- [Assigning customized privileges to VMwareDisks agent](#)

About the roles and privileges assigned in vCenter

The following set of privileges are available after you install the Veritas High Availability Console. These privileges define the operations that a user can perform on the system. You can create roles and then assign privileges to them or assign privileges to the existing roles that are available in the vSphere environment. Application monitoring operations are enabled or disabled depending on the privileges that are assigned to the vCenter user account. For example, the Admin privilege is required for configuring application monitoring on a system.

vCenter Server administrators can use these privileges to configure access control while monitoring an application.

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

Can perform all 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 maintenance mode, and view application status.

Assigning customized privileges to VMwareDisks agent

This section describes a procedure exclusively for users who want to assign customized privileges to the VMwareDisks agent on an ESX host, instead of administrative or root user privileges. If you want to assign administrative or root user privileges, skip this section.

The following subsections describe the background and the configuration workflow in detail.

- [About assigning privileges to VMwareDisks agent](#)
- [Creating a role with customized privileges for VMwareDisks agent](#)
- [Creating an ESX user account](#)
- [Integrating an ESX user account with Active Directory](#)
- [Assigning a role to an ESX user account](#)

About assigning privileges to VMwareDisks agent

The application monitoring configuration for Cluster Server (VCS) agents in a VMware virtual environment involves the VMwareDisks agent. In the event of an application failure, the VMwareDisks agent sends a disk-detach request to the ESX host, and then attaches the disk to the failover target system.

To enable the VMwareDisks agent to communicate with the ESX host, during the application monitoring configuration workflow, you must specify an ESX user account. The specified ESX user account must have administrative privileges, or should be a root user. If the ESX user account does not have these privileges, you must create a role, add certain privileges to the created role, and then assign the role to the ESX user account.

If you do not want to assign the role to an existing ESX user account, you can create a new ESX user account, and then assign the role. You can further integrate the new ESX user account with an Active Directory-based authentication service if available in the VMware environment. The VMwareDisks agent can then use the same user account to perform its tasks on all ESX hosts linked to the Active Directory.

Creating a role with customized privileges for VMwareDisks agent

This section provides the steps to create a role with adequate privileges that the VMwareDisks agent can use in an ESX cluster. The assigned privileges do not include administrative or root user privileges:

To create a role

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Administration > Roles**.
- 2 Click **Add Role**.
- 3 On the Add New Role panel, specify a name for the new role. For example, "VMwareDisks Resources Manager".
- 4 To select privileges for the VMwareDisks Resources Manager role, in the Privileges tree, click the following check boxes.
 - **All Privileges > Datastore > Low level file operations**.
 - **All Privileges > Virtual Machine > Configuration > Adding existing disk**.
 - **All Privileges > Virtual Machine > Change resource**.
 - **All Privileges > Virtual Machine > Configuration > Remove disk**.
- 5 Click **OK**.

Creating an ESX user account

If you want to assign the role that you created in the section [Creating a role with customized privileges for VMwareDisks agent](#), to an existing user, skip this section. Proceed to the section [Assigning a role to an ESX user account](#).

If you want to create a new user account to assign the new role to, perform the steps described in this section:

To add a local ESX user

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 2 In the left pane, click the ESX host and in the right pane, click **Local Users & Groups**.

The Users list appears by default.
- 3 If the Users list is not displayed, to view the Users list, on the View bar, click **Users**.
- 4 To add a new user, right-click any existing user, and click **Add**.

- 5 In the **Add New User** panel, specify a Login and Password to define a new user account for configuring VMwareDisks resources.

To confirm the password, retype the password.

To define the new user account, you can also specify a descriptive User Name and user ID (UID). If you do not specify the UID, the vCenter server automatically assigns one.

- 6 Click **OK**.

Integrating an ESX user account with Active Directory

After you create a new ESX user account for the VMwareDisks agent to communicate with an ESX host, you can optionally integrate the account with any existing Active Directory authentication in your environment. Else, the new ESX user account depends on the local authentication mechanism on the ESX host, and you will need to configure one account per host.

Integrating with an existing Active Directory mechanism helps you leverage the same ESX user account across multiple ESX hosts for VMwareDisks agent configurations.

To integrate with Active Directory

- 1 Create a domain user in the Active Directory.
- 2 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 3 In the left pane, click the ESX host and in the right pane, click **Configuration**.
- 4 In the Software panel, click **Authentication Services**.
- 5 Review the Directory Services Configuration.
If the Directory Service Type is not Active Directory, in the top right corner, click **Properties**.
- 6 In the Directory Service Configuration panel, from the Select Directory Service Type drop down list, select **Active Directory**.
- 7 In the Domain Settings area, specify the **Domain**, and click **Join Domain**.
Alternatively, configure vSphere Authentication proxy.
- 8 Enter the user name and password of a directory service user that has permissions to join the host to the domain, and click **OK**.

Assigning a role to an ESX user account

This section describes the steps to assign a role to an ESX user:

To assign the role to a user

- 1** Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 2** In the left pane, click the ESX host and in the right pane, click **Permissions**.
- 3** In the Permissions tab, right-click the blank space, and click **Add Permission**.
- 4** In the Assign Permissions panel, click **Add**.
- 5** In the Users and Groups frame of the Select Users and Groups panel, specify the user(s) that you want to assign the new role for storage management (for example, VMwareDisks Resources Manager).

Press the **Ctrl** key and click to select multiple users, if required, and then click **Add** and click **OK**.
- 6** In the Assigned Role drop down list, click the new role (for example, VMwareDisks Resources Manager), and then click **OK**.

Troubleshooting

This appendix includes the following topics:

- [Agent logging on virtual machine](#)
- [Troubleshooting wizard-based configuration issues](#)
- [Troubleshooting issues with the Veritas High Availability view](#)

Agent logging on virtual machine

Veritas High Availability agents generate log files that are appended by letters. Letter A indicates the first log file, B the second, C the third, and so on.

The agent log components are defined as follows:

- **Timestamp:** the date and time the message was generated.
- **Mnemonic:** the string ID that represents the product (for example, VCS).
- **Severity:** levels include CRITICAL, ERROR, WARNING, NOTICE, and INFO (most to least severe, respectively).
- **UMI:** a unique message ID.
- **Message Text:** the actual message generated by the agent.

The agent logs are located in the following location:

```
/var/VRTSvcs/log/<agentname>_A.log
```

The format of the agent log is as follows:

```
Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI | Agent Type |  
Resource Name | Entry point | Message text
```

A typical agent log resembles:

```
2012/08/15 13:34:44 VCS ERROR V-16-2-13067 Thread(4146068336) Agent
is calling clean for resource(MQ1) because the resource became OFFLINE
unexpectedly, on its own.
```

Troubleshooting wizard-based configuration issues

This section lists common troubleshooting issues that you may encounter during or after the steps related to the VCS Cluster Configuration Wizard and the Veritas High Availability Configuration Wizard.

Veritas High Availability Configuration wizard displays the "hadiscover is not recognized as an internal or external command" error

While configuring application monitoring the Veritas High Availability Configuration wizard may display the "hadiscover is not recognized as an internal or external command" error, after you click Next on the Application Selection panel.

This issue occurs if you launch the wizard from a system where you have reinstalled VCS.

Workaround:

Exit the wizard, restart the Veritas Storage Foundation Messaging Service and then re-run the wizard.

Running the 'hastop -all' command detaches virtual disks

The `hastop -all` command takes offline all the components and components groups of a configured application, and then stops the VCS cluster. In the process, the command detaches the virtual disks from the VCS cluster nodes. (2920101)

Workaround: If you want to stop the VCS cluster (and not the applications running on cluster nodes), instead of the "hastop -all", use the following command:

```
hastop -all -force
```

This command stops the cluster without affecting the virtual disks attached to the VCS cluster nodes.

Validation may fail when you add a failover system

On the Configuration Inputs panel, when you add a failover system using the Add System option, you may see the following error message:

Validation has failed for <System Name>.

Workaround: Verify the following:

- The Veritas High Availability guest components are installed on the system.
- The operating system running on the system is supported by this product.
- The specified system name or IP address is valid and the system is switched on and accessible over the network.
- The firewall settings allow access on port 5634 used by the Storage Foundation Messaging Service.
- If xprtd daemon is running on the system, restarting the xprtd service may resolve the issue.

Adding a failover system may fail if you configure a cluster with communication links over UDP

When you configure a cluster with communication links over UDP and add a failover system, you may see the following error message:

The same network adapter is specified for one or more links on the system. You must select a different network adapter for each communication link.

This issue occurs when you unplumb an IP address from one of the configured communication links. In this scenario, VCS may also go into jeopardy.

Workaround:

Plumb the IP address again.

Troubleshooting issues with the Veritas High Availability view

This section lists common troubleshooting scenarios that you may encounter when using the Veritas High Availability view.

Veritas high availability view is not visible from a cluster system

The Veritas High Availability view displays the cluster view (consolidated cluster-level health information of the configured application/s running on the selected system). In some multi-node clusters, the view is not visible from at least one of the cluster system.

This behavior occurs if connectivity of the configured LLT links fail. This may be a networking error. (2863649)

Workaround

Ensure that valid LLT links are configured for the affected cluster system, and then retry.

Veritas High Availability view does not display the application monitoring status

The Veritas High Availability view 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 Veritas High Availability view:

- Verify that the Veritas High Availability Console host is running and is accessible over the network.
- Verify that the VMware Web Service is running on the vCenter Server.
- Verify that the VMware Tools Service is running on the guest virtual machine.
- Verify that the Veritas Storage Foundation Messaging Service (xpirtld process) is running on the Veritas High Availability Console and the virtual machine. If it is stopped, type the following on the command prompt:

```
net start xpirtld
```
- Verify that ports 14152, 14153, and 5634 are not blocked by a firewall.
- Log out of the vSphere Client and then login again. Then, verify that the Veritas High Availability plugin is installed and enabled.

Veritas High Availability view may freeze due to special characters in application display name

For a monitored application, if you specify a display name that contains special characters, one or both of the following symptoms may occur:

- The Veritas high availability view may freeze
- The Veritas high availability view may display an Adobe exception error message. Based on your browser settings, the Adobe exception message may or may not appear. However, in both cases the tab may freeze. (2923079)

Workaround:

Reset the display name using only those characters that belong to the following list:

- any alphanumeric character

- space
- underscore

Use the following command to reset the display name:

```
hagrp -modify sg name UserAssoc -update Name "modified display  
name without special characters"
```

If the Console host abruptly restarts, the high availability view may disappear

If the system that hosts the Veritas High Availability Console abruptly restarts, then the registration of the Veritas High Availability plugin with the VMware vCenter Server does not persist.

As a result, the Veritas High Availability view does not appear on the vSphere Client GUI. (2919549)

Workaround:

Repair the Console installation.

Veritas high availability view may fail to load or refresh

The Veritas High Availability view displays health information of monitored applications in a VCS cluster. The tab display may fail to load. It may alternatively fail to refresh itself after the default interval of 60 seconds. (2932028)

Workaround:

Restarting the xprtd service may resolve the issue.

To restart the xprtld service

1 Stop the xprtld service:

For RHEL 7, SLES 12, and supported RHEL distributions:

```
# systemctl stop xprtld
```

For earlier versions of RHEL, SLES, and supported RHEL distributions:

```
# /etc/init.d/xprtld stop
```

2 Ensure that xprtld is stopped:

```
# ps -ef | grep xprtld
```

If the services is not stopped, terminate the process:

```
# kill -9 xprtld_pid_value
```

Where pid is the process ID of the xprtld process.

3 Start xprtld service:

For RHEL 7, SLES 12, and supported RHEL distributions:

```
# systemctl start xprtld
```

For earlier versions of RHEL, SLES, and supported RHEL distributions:

```
# /etc/init.d/xprtld start
```

Operating system commands to unmount resource may fail

If a user configures a mount point as component (resource) while configuring high availability for an application, then unconfiguring the application from the Veritas High Availability view may lock the mount point. Operating system commands to unmounts the resource may fail. (3574657)

Workaround

Ensure that you stop the application before you unconfigure the application by clicking **More > Unconfigure Application Monitoring** in the Veritas High Availability view.