

Cluster Server 7.3.1 Configuration Guide for Custom Applications - Windows

Last updated: 2017-11-05

Legal Notice

Copyright © 2017 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
500 E Middlefield Road
Mountain View, CA 94043

<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:

doc.feedback@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	6
	How does the Veritas High Availability solution work	6
	How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host	8
	Typical VCS cluster configuration in a virtual environment	9
	How does the Veritas High Availability solution integrate with vSphere Web Client	10
	Supported VMware versions	12
	Managing storage	12
 Chapter 2	 Configuring application monitoring using the Veritas High Availability solution	 14
	Deploying the Veritas High Availability solution for configuring application monitoring	14
	Notes and recommendations	16
	Assigning privileges for non-administrator ESX/ESXi user account	18
	Configuring application monitoring	21
	Configuring the VCS cluster	22
	Configuring the application	25
	Modifying the ESXDetails attribute	31
 Chapter 3	 Administering application monitoring	 33
	About the various interfaces available for performing application monitoring tasks	33
	Administering application monitoring using the Veritas High Availability tab	34
	Understanding the Veritas High Availability tab work area	35
	To configure or unconfigure application monitoring	38
	To start or stop applications	39
	To switch an application to another system	40
	To add or remove a failover system	41
	To suspend or resume application monitoring	45

To clear Fault state	46
To resolve a held-up operation	47
To determine application state	47
To remove all monitoring configurations	47
To remove VCS cluster configurations	48
Administering application monitoring settings	48
Administering application availability using Veritas High Availability	
dashboard	49
Understanding the dashboard work area	50
Monitoring applications across a data center	54
Monitoring applications across an ESX cluster	54
Searching for application instances by using filters	54
Selecting multiple applications for batch operations	55
Starting an application using the dashboard	55
Stopping an application by using the dashboard	56
Entering an application into maintenance mode	56
Bringing an application out of maintenance mode	57
Switching an application	57
Resolving dashboard alerts	58

Appendix A	Troubleshooting	60
	Troubleshooting application monitoring 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	61
	Adding a failover system may fail if you configure a cluster with	
	communication links over UDP	61
	Troubleshooting Veritas High Availability view issues	62
	Veritas High Availability tab not visible from a cluster node	62
	Veritas High Availability tab does not display the application	
	monitoring status	62
	Veritas High Availability tab may freeze due to special characters	
	in application display name	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 does the Veritas High Availability solution work](#)
- [Supported VMware versions](#)
- [Managing storage](#)

How does the Veritas High Availability solution work

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 8.

In 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 fail over to the failover target system. During the fail over, 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.

In event of a failure in a site recovery configuration, the following tasks are performed for application monitoring continuity:

- The virtual machines at the protected site are failed over to the recovery site.
- The pre-online script defined in the form of a command in the SRM recovery plan applies the specified attribute values for the application components.
- The status monitoring script retrieves the application status.
- The network agents bring the network components online and the application-specific agents start the application services on the failover target system.

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 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.

See [“Modifying the ESXDetails attribute”](#) on page 31.

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.

Typical VCS cluster configuration in a virtual environment

A typical VCS cluster configuration in a VMware virtual environment involves two or more virtual machines. The virtual machine on which the application is active, accesses a non-shared VMware VMDK or RDM disk that resides on a VMware datastore.

The virtual machines involved in the VCS cluster configuration may belong to a single ESX host or could reside on separate ESX hosts. If the virtual machines reside on separate ESX hosts, the datastore on which the VMware VMDK or RDM disks (on which the application data is stored) reside must be accessible to each of these ESX hosts.


The application binaries are installed on the virtual machines and the data files are installed on the VMware disk drive. The VCS agents monitor the application components and services, and the storage and network components that the application uses.

During a failover, the VCS storage agents (MountV-VMNSDg-VMwareDisks in case of SFW storage, Mount-NativeDisks-VMwareDisks in case of LDM storage) move the VMware disks to the new system. The VCS network agents bring the network components online, and the application-specific agents then start the application services on the new system.

In a site recovery environment, Veritas High Availability solution additionally provides script files for the following tasks. These files are invoked when the SRM recovery plan is executed.

- Set up communication between the vCenter Server and the SRM Server at the recovery site and the virtual machines at the protected site.
- Assign a SiteID to both the sites.
- Specify attribute values for the application components at the respective site.
- Retrieve the application status in the SRM recovery report, after the virtual machine is started at the recovery site.


Figure 1-1 Typical cluster configuration in a VMware virtual environment



How does the Veritas High Availability solution integrate with vSphere Web Client

The Veritas High Availability solution integrates with a vSphere Web Client to enable you to perform the application monitoring tasks from within a vSphere Web Client interface.

The following diagram illustrates how the Veritas High Availability solution integrates with vSphere Web Client:



The Veritas High Availability solution provides a Veritas HA Plug-in for integration with VMware vSphere Web Client, and adds the following interfaces for performing the application monitoring tasks:

- **Veritas High Availability tab:** From this view you configure application monitoring and then monitor and control the configured application on that virtual machine. After configuring application monitoring, the Veritas High Availability tab view displays the state of the configured application.
- **Veritas High Availability dashboard:** This view appears at an ESX cluster or datacenter-level. The Veritas High Availability Dashboard displays aggregate health information for all the monitored applications running in that cluster or datacenter.

These interfaces appear in the Monitor tab of vSphere Web Client, after you install the Veritas HA Plug-in for vSphere Web Client add-on for the Veritas InfoScale Operations Manager (VIOM) Management Server. You must install this add-on on a VIOM Management Server and register the Veritas HA Plug-in for vSphere Web Client with all the required vCenter servers.

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 and displays the information in the Veritas High Availability tab and the dashboard.

Supported VMware versions

For the latest list of supported VMware Servers and Management Clients refer to the software compatibility list (SCL) at:

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

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 volumes on the virtual disks.

Note: If your storage configuration involves NetApp filers that are directly connected to the systems using iSCSI initiator, you cannot configure application monitoring in a virtual environment with non-shared disks.

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

- If the storage is managed using SFW, the MountV, VMNSDg, and VMwareDisks agents are used.
- If the storage is managed using LDM, the Mount, NativeDisks, and VMwareDisks 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*.

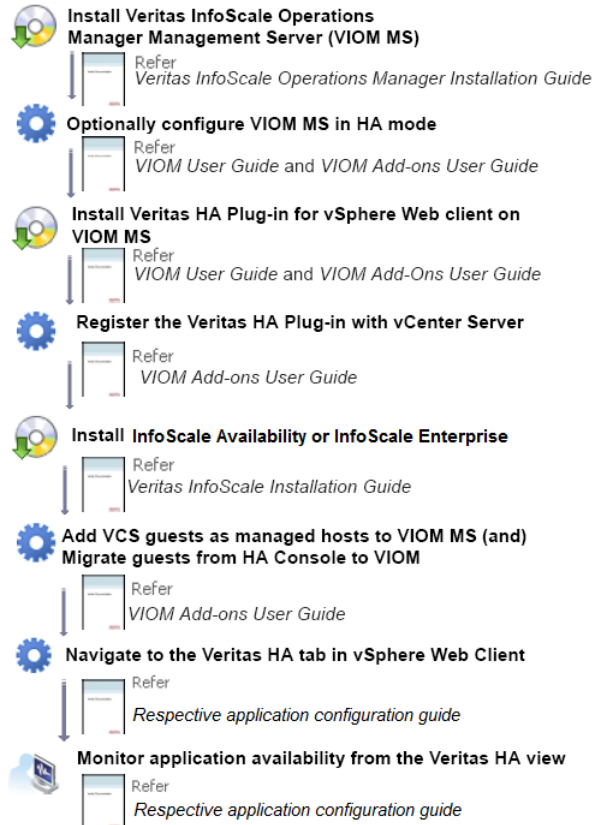
Configuring application monitoring using the Veritas High Availability solution

This chapter includes the following topics:

- [Deploying the Veritas High Availability solution for configuring application monitoring](#)
- [Notes and recommendations](#)
- [Configuring application monitoring](#)
- [Modifying the ESXDetails attribute](#)

Deploying the Veritas High Availability solution for configuring application monitoring

The following figure depicts the workflow for deploying the Veritas High Availability solution for configuring application monitoring.



Notes:

- You must follow this workflow only if you plan to configure application monitoring using the vSphere Web Client infrastructure. Alternatively, you can configure application monitoring and perform the further administrative tasks directly from a browser window using the following URL:
https://<virtualmachineNameorIPaddress>:5634/vcs/admin/application_health.html?priv=ADMIN
- For the configuring application monitoring in a VMware SRM environment, see:
<http://www.symantec.com/docs/DOC7958>
- If you plan to migrate from an existing Veritas High Availability vSphere Desktop Client integrated environment to a vSphere Web Client integrated environment, see:
https://sort.symantec.com/public/vom/Symantec_HA_Plug-in_for_vSphere_Web_Client_README.pdf

Notes and recommendations

Note the following prerequisites before configuring application monitoring:

- Verify that the boot sequence of the virtual machine is such that the boot disk (OS hard disk) is placed before the removable disks.
If the sequence places the removable disks before the boot disk, the virtual machine may not reboot after an application failover. The reboot may halt with an "OS not found" error.
This issue occurs because during the application failover the removable disks are detached from the current virtual machine and are attached on the failover target system.
- Verify that VMware Tools is installed on the virtual machine.
Install the version that is similar to or later than that available with VMware ESX 4.1.
- Verify that all the systems on which you want to configure application monitoring belong to the same domain.
- Verify that the ESX/ESXi host user account has administrative privileges or is a root user.
If the ESX/ESXi user account fails to have the administrative privileges or is not a root user, then in event of a failure the disk deattach and attach operation may fail.
If you do not want to use the administrator user account or the root user, then you must create a role, add the required privileges to the created role and then add the ESX user to that role.
See ["Assigning privileges for non-administrator ESX/ESXi user account"](#) on page 18.
- If you want to monitor the mount points, ensure that the volumes or mounts selected during configuration do not have multiple paths or drive letters assigned to it.
The Veritas High Availability Configuration wizard fails to configure monitoring if the selected mount points or a volume has multiple drive letters or paths are assigned to it.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Veritas High Availability installer, wizard, and services.
- After configuring services, processes, and mount points for monitoring, if you create another service, process, or mount point, then these new components are not monitored as part of the existing configuration.

In this case, you can either use the VCS commands to add the components to the configuration or unconfigure the existing configuration and then run the wizard again to configure all the components.

- In case the VMwareDisks agent resource is configured manually, care should be taken not to add the operating system disk in the configuration. The VMwareDisks agent does not block this operation. This might lead to a system crash during failover.
- If VMware vMotion is triggered at the same time as an application fails over, the VMwareDisks resource may either fail to go offline or may report an unknown status. The resource will eventually failover and report online after the vMotion is successful and the application is online on the target system.
- VMware snapshot operations may fail if VMwareDisks agent is configured for a physical RDM type of disk. Currently only virtual RDM disks are supported.
- Non-shared disks partitioned using GUID Partition Table (GPT) are not supported. Currently only Master Boot Record (MBR) partition is supported.
- VMwareDisks agent does not support disks attached to the virtual machine using IDE controllers. The agent resource reports an unknown if IDE type of disks are configured.
- In case VMware HA is disabled and the ESX itself faults, VCS moves the application to the target failover system on another ESX host. VMwareDisks agent registers the faulted system on the new ESX host. When you try to power on the faulted system, you may see the following message in the vSphere Client:

```
This virtual machine might have been moved or copied.  
In order to configure certain management and networking features,  
VMware ESX needs to know if this virtual machine was moved or copied.  
If you don't know, answer "I copied it".
```

You must select “I moved it” (instead of the default “I copied it”) on this message prompt.

- You must not restore a snapshot on a virtual machine where an application is currently online, if the snapshot was taken when the application was offline on that virtual machine. Doing this may cause an unwanted fail over.
This also applies in the reverse scenario; you should not restore a snapshot where the application was online on a virtual machine, where the application is currently offline. This may lead to a misconfiguration where the application is online on multiple systems simultaneously.
- If you want to suspend a system on which an application is currently online, then you must first switch the application to a failover target system.

If you suspend the system without switching the application, then VCS moves the disks along with the application to another system.

Later, when you try to restore the suspended system, VMware does not allow the operation because the disks that were attached before the system was suspended are no longer with the system.

- While creating a VCS cluster in a virtual environment, you must configure one of the cluster communication link over a public adapter in addition to the link configured over a private adapter. To have less VCS cluster communication over the link using the public adapter, you may assign it low priority. This keeps the VCS cluster communication intact even if the private network adapters fail. If the cluster communication is configured over the private adapters only, the cluster systems may fail to communicate with each other in case of network failure. In this scenario, each system considers that the other system has faulted, and then try to gain access to the disks, thereby leading to an application fault.
- VMware Fault Tolerance does not support adding or removing of non-shared disks between virtual machines. During a failover, disks that contain application data cannot be moved to alternate failover systems. Applications that are being monitored thus cannot be brought online on the failover systems.
- For cluster communication, you must not select the teamed network adapter or the independently listed adapters that are a part of the teamed NIC. A teamed network adapter is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address, due to which you may experience the following issues:
 - The application monitoring configuration wizard may fail to discover the specified network adapters
 - The application monitoring configuration wizard may fail to discover/validate the specified system name

Assigning privileges for non-administrator ESX/ESXi user account

The application monitoring configuration in a VMware virtual environment using non-shared disks involves the VMwareDisks agent. In event of a failure, the VMwareDisks agent sends a disk detach request to the ESX/ESXi host and then attaches it to the new failover target system.

To enable the VMwareDisks agent to communicate with the ESX/ESXi host, we need to specify the ESX user account details during the application configuration workflow. This ESX user account must have the administrative privileges or should be a root user. If the ESX user account does not have these privileges, you must perform the following tasks:

- Create a role having the following privileges

- Low level file operations
 - Add existing disk
 - Change resource
 - Remove disk
- See [“Creating a role”](#) on page 19.
- Integrate with the existing authentication mechanism
 See [“Integrating with Active Directory or local authentication”](#) on page 20.
 - Add the ESX user to the created role
 See [“Adding a user to the role”](#) on page 21.

Note: If you do not want to add the existing user, you can create a new user and then add the same to the created role

See [“Creating a new user”](#) on page 20.

Creating a role

Perform the following steps to create the 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, "ESX/ESXi User Role for Application Monitoring".
- 4 In the Privileges tree, click the following check boxes to assign the required privileges:
 - **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**.

Integrating with Active Directory or local authentication

To integrate with Active Directory or local authentication

- 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 Click the ESX host.
- 4 In the right pane, click **Configuration**.
- 5 In the Software panel, click **Authentication Services**.
- 6 Review the Directory Services Configuration.

If the Directory Service Type is not Active Directory, and you do not want to integrate with Active Directory, proceed to the section,

See [“Adding a user to the role”](#) on page 21.

If the Directory Service Type is not Active Directory, and you want to integrate with Active Directory, in the top right corner, click **Properties**.

- 7 In the Directory Service Configuration panel, from the Select Directory Service Type drop down list, select **Active Directory**.
- 8 In the Domain Settings area, specify the **Domain**, and click **Join Domain**.
Alternatively, configure vSphere Authentication proxy.
- 9 Enter the user name and password of a directory service user that has permissions to join the host to the domain, and click **OK**.

Creating a new user

You must perform this task only if you do not want to add the existing user to the created role.

Perform the following steps to create a new user

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 2 Click the ESX host.
- 3 In the right pane, click **Local Users & Groups**.
The Users list appears by default.
- 4 If the Users list is not displayed, on the View bar, click **Users**.
Alternatively, if the Users list is displayed, right-click any existing user and then click **Add**.

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

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**.

Adding a user to the role

To add a user to the role

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 2 Click the ESX host.
- 3 In the right pane, click **Permissions**.
- 4 In the Permissions tab, right-click the blank space, and click **Add Permission**.
- 5 In the Assign Permissions panel, click **Add**.
- 6 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.

Press Ctrl and click to select multiple users, if required, and then click **Add** and click **OK**.
- 7 In the Assigned Role drop down list, click the new role and then click **OK**.

Configuring application monitoring

Configuring an application for monitoring involves the following tasks:

1. Configuring the VCS cluster

This task involves selecting the virtual machines on which you want to configure monitoring and setting up the network communication links between those virtual machines.

See [“Configuring the VCS cluster”](#) on page 22.

2. Configuring the application

This task involves configuring the application in one of the following modes:

- Start/stop mode on a single system: In the event of a failure, the application is restarted on a virtual machine for a configured number of times.

- **Failover mode on multiple systems:** In the event of a failure, the application is first restarted on a virtual machine for a configured number of times. If the application does not restart, the virtual machine fails over to another ESX host.

When you configure an application in the failover mode, you need to select the ESX hosts where the virtual machines can fail over.

See [“Configuring the application”](#) on page 25.

Configuring the VCS cluster

VCS cluster configuration involves selecting the virtual machines on which you want to configure monitoring and setting up the network communication links between the selected virtual machines.

To configure the VCS cluster

- 1** Launch the vSphere Client and connect to the vCenter Server that manages the virtual machine. On the Inventory view of the vCenter Server, in the left pane, select the virtual machine where you want to configure application monitoring. Then, in the right pane, select the Veritas High Availability tab.
- 2** On the Veritas High Availability view, click **Configure a VCS Cluster**. This launches the VCS cluster configuration wizard. Unless you configure a cluster, you cannot configure an application for monitoring.
- 3** On the Welcome panel, review the pre-requisites and then click **Next**.

- On the Configuration Inputs panel, specify the systems for the VCS cluster operations and the user account details for each system.

Note: The specified user account must have administrative privileges on the system.

The **Cluster systems** lists the systems that are included in the cluster configuration. The local system is selected by default.

To add more systems, click **Add System** and then on the Add System dialog box, specify the following details of the system that you want to add to the VCS cluster.

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. The user name must be in the domain.com\username. Note: The specified user must be a domain user having administrative privileges on all the selected system.
Password	Specify the password for the user account mentioned.
Use the specified user account on all systems	Uses the specified user account on all the cluster systems. This option is selected by default.

The wizard validates the system details and then adds the system to VCS cluster system list.

- Skip this step if you do not want to modify the default security settings for your cluster.

To modify the security settings for the cluster, on the Configuration Inputs panel, click **Advanced Settings** . In the Advanced settings dialog box, specify the following details and then 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.

- 6 On the Network Details panel, select the type of communication for the VCS cluster and then select the adapters to configure the communication links.

Depending on the network over which you want to configure the links, select:

- **Use MAC address for cluster communication (LLT over Ethernet) :**
The LLT over Ethernet communication configures the links over the non-routed network. Choose this mode only if the failover target systems reside in the same subnet.
- **Use IP address for cluster communication (LLT over UDP):** The LLT over UDP communication configures the links over the routed network. You choose this mode regardless of whether the failover target systems reside in the same subnet or in different subnets. You can select only those adapters that have an IP address.

You must select a minimum of two adapters per system.

Veritas recommends the following:

- IP addresses that are assigned to the selected adapters should be in different subnets.
- One of the network adapters must be a public adapter. You may assign low priority to the VCS cluster communication link that uses the public adapter.

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>To configure links over UDP, select the type of IP protocol and then specify the required details for each communication link.</p> <p>Note: Do not select the teamed network adapters and the independently listed adapters that are a part of the teamed NIC.</p>
IP Address	Specify the IP address 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>A specified port for a link is used for all the cluster systems on that link.</p>
Subnet mask	Displays the subnet masks to which the specified IP belongs.

By default, the VCS cluster communication link that uses the public adapter is configured as low-priority link. To change the priority, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

For information about adding or changing the selected network links after the configuration workflow is complete, see the *Cluster Server Administrator's Guide*.

- 7 On the Configuration Summary panel, specify a cluster name and a cluster ID of your choice. Review the VCS cluster configuration details and then click **Next** to initiate the configuration.

If the network contains multiple clusters, the wizard verifies the cluster ID with the IDs assigned to all the accessible clusters in the network. The wizard does not validate the assigned ID with the clusters that are not accessible during the validation. Veritas recommends you to validate the uniqueness of the assigned ID in the existing network.

- 8 On the Implementation panel, the wizard displays the VCS cluster configuration tasks and the status of each task. After all the tasks are complete, click **Next**.

If the configuration tasks fail, 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.

- 9 On the Finish panel, click **Finish** to complete the wizard workflow. This step completes the VCS cluster configuration.

The Veritas High Availability view now displays the link to configure an application for high availability.

Configuring the application

Perform the following steps to configure monitoring for services, processes, and mount points using the Veritas High Availability Configuration Wizard.

Note: You can configure monitoring for multiple services and processes in a single wizard workflow. However, you cannot configure multiple applications simultaneously. To configure another application, run the wizard again.

To configure the application

- 1 Launch the vSphere Client and connect to the vCenter Server that manages the virtual machine. On the Inventory view of the vCenter Server, in the left pane, select the virtual machine where you want to configure application monitoring. Then, in the right pane, select the Veritas High Availability tab.
- 2 On the Veritas High Availability tab, click **Configure Application for High Availability**. This option is available only after a cluster has been configured. Unless you configure a cluster, you cannot configure an application for monitoring.
- 3 On the Welcome panel of the Veritas High Availability Configuration wizard, review the pre-requisites and then click **Next**.
- 4 On the Application Selection panel, select the application from the Supported Applications list and then click **Next**.

Alternatively, you can use the **Search** box to find the application and then click **Next**.

If you want to download any of the High Availability Agents, click the **Download Application Agents (SORT)** link to download the agents from the Veritas Operations Readiness Tools (SORT) site.

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

- 5 On the Windows Service Selection panel, select the services that you want to monitor and then click **Next**.

Alternatively, you can use the **Search** box to find the service and then click **Next**.

The wizard automatically discovers the services on the system. If a selected service depends on some other services, you must also select those services. You can define the dependencies between those services on the Define Start-Stop Order panel later.

If you do not want to monitor any services, click **Next** without selecting any service.

- 6 On the Windows Process Selection panel, specify the processes that you want to monitor.

To specify a process, click **Add Process**.

On the Process Parameters dialog box, select the type of monitoring that you want to configure, specify the required details, and then click **OK**.

For direct process monitoring, specify the following details:

Process Full Path	Specify the complete path of the process executable file including its extension.
Arguments	Specify the command line arguments, if any, for the process to be monitored.
Run process in local system account context	By default the specified process runs in the context of the local system account.
Run process in specified user account context	Select to run the process in a different user's context, and then specify the user name and password in the respective fields. The user name must be in the format user@domain.com or domain.com\username. Note: Ensure that you specify a valid user name and that the user account has adequate privileges on the system where you want to configure application monitoring. Otherwise, application monitoring may fail.

For program-based process monitoring, specify the following details:

Start program	Specify the full path of the program that starts the process to be configured for monitoring.
Arguments	Specify the command line arguments, if any, for the start program.
Monitor program	Specify the full path of the program that monitors the process to be configured for monitoring.
Arguments	Specify the command line arguments, if any, for the monitor program.
Stop program	Specify the full path of the program that stops the process to be configured for monitoring.
Arguments	Specify the command line arguments, if any, for the stop program.

Run process in local system account context	By default the specified process runs in the context of the local system account.
Run process in specified user account context	<p>Select to run the process in a different user's context, and then specify the user name and password in the respective fields.</p> <p>The user name must be in the format user@domain.com or domain.com\username.</p> <p>Note: Ensure that you specify a valid user name and that the user account has adequate privileges on the system where you want to configure application monitoring. Otherwise, application monitoring may fail.</p>

The process or the programs that you add is displayed on the Windows Process Selection panel. Repeat this step for all the processes that you want to configure for monitoring.

If you do not want to monitor any processes, click **Next** without specifying any process.

- 7 On the Mount Point Selection panel, select the mount points that you want to monitor and then click **Next**.

If you do not want to monitor any mount points, click **Next** without selecting any mount points.

- 8 On the Define Start-Stop Order panel, specify the order in which you want the selected services, processes, and mount points to be started or stopped.

To define the dependency between the components, select an application component from the **Parent Component** box and then select the components from the **Depends on** box.

While starting the service or process, the components are brought online in the defined order. For example, if a service is dependent on a mount point, then while starting the service, the mount point is first brought online and then the service itself.

Click **Next** to proceed with the configuration.

- On the Configuration Inputs panel, select the VCS cluster systems on which you want to configure the application for high availability and move them to the **Application failover targets** list. The local system is selected by default.

Using the up-down arrow keys, you can define the priority order for the failover systems.

For each system that you assign as a failover target, you must specify the domain user account details in the appropriate fields on the Edit System dialog box. The VCS agents use these details to perform domain operations (such as Active Directory updates).

- On the Virtual Network Details panel, specify the virtual IP and the network details for the application to be configured and then click **Next**.

To specify the virtual IP and network details, select the IP protocol and then specify the following details for each failover system:

Note: You must select the same IP protocol as the one that was selected during the VCS cluster configuration.

Virtual IP address	Specify a unique virtual IP address.
Subnet mask	Specify the subnet mask to which the IP address belongs.
Virtual name	Specify a virtual name.
Network Adapter column	Select the network adapter that will host the virtual IP.

- 11 On the Storage HA Inputs panel, specify the ESX/ESXi hosts and the administrative user account details for each host, and then click **Next**.

Note: This panel appears only if you have specified a mount point, or multiple application failover targets, or both for the selected custom application.

To specify the ESX/ESXi hosts, click **Add ESX/ESXi Host** and on the Add ESX/ESXi Host dialogue box, specify the following details:

ESX/ESXi hostname or IP address	<p>Specify the target ESX hostname or IP address.</p> <p>The virtual machines will fail over on this ESX host during vMotion.</p> <p>The mount points configured on the ESX host where the application is currently running must be available on the target ESX host.</p>
User name	<p>Specify a user account for the ESX host.</p> <p>The user account must have administrator privileges on the specified ESX host.</p>
Password	<p>Specify the password for the user account provided in the User name text box.</p>

Note: By default, the wizard sets up a communication link with the ESX/ESXi server. You can modify the configuration to set up the communication link with vCenter Server instead. To set up a link with a vCenter Server, you must modify the ESXDetails attribute after this application monitoring configuration workflow is complete.

See [“Modifying the ESXDetails attribute”](#) on page 31.

- 12 On the Configuration Summary panel, review the application configuration details.

The wizard assigns a unique name to the application service group. Click **Rename** to rename the service group.

Click **Next** to initiate the application monitoring configuration.

- 13 On the Implementation panel, the wizard performs the application configuration tasks.

The wizard displays the status of each task. After all the tasks are complete, click **Next**.

If the configuration tasks fail, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure, and run the wizard again to configure application monitoring.

- 14 On the Finish panel, click **Finish** to complete the wizard workflow. This configures the application for high availability on the selected VCS cluster systems.

The Veritas High Availability view now displays the application status and its component dependency.

If the application status shows as not running, click **Start** to start the configured components on the system.

Modifying the ESXDetails attribute

You must modify the value of the "ESXDetails" attribute (of the VMwareDisks agent) if you want the VMwareDisks agent to communicate with the vCenter Server (instead of the ESX/ESXi host) for the disk detach and attach operations.

By default the "ESX Details" attribute of the VMwareDisks agent used the host names or IP addresses and the user account details of the ESX hosts on which the virtual machines are configured. To enable the VMwareDisks agent to communicate with the vCenter Server, you must modify the ESXDetails attribute and provide the host name or IP address and the user account details of the vCenter Server to which the virtual machines belong.

Use the Cluster Manager (Java Console) or the Command Line to modify the attribute values.

To modify the attribute from Cluster Manager

- 1 From the Cluster Manager configuration tree, select the VMwareDisks resource and then select the **Properties** tab.
- 2 On the **Properties** tab, click the Edit icon next to the ESX Details attribute.
- 3 On the Edit Attribute dialog box, select all the entries specified under the Key-Value column and press "-" to delete them.
- 4 Encrypt the password of the vCenter Server user account.

For details on encrypting agent passwords, see the *Cluster Server Administrator's Guide*.

- 5 On the Edit Attribute dialog box, click “+” to specify the values under the Key-Value column.
- 6 Under the Key column, specify the vCenter Server host name or the IP address.
- 7 Under the Value column, specify the encrypted password of the vCenter Server user account (from step 4)
- 8 Click **Ok** to confirm the changes.
- 9 Repeat the steps for all VMwareDisks resources from the Cluster Manager configuration tree.
- 10 Save and close the configuration.

To modify/specify the attribute from Command Line

- 1 Change the VCS configuration to read/write mode.

```
Haconf -makerw
```

- 2 Delete the existing details of the ESX Server.

```
hares -modify VMwareDisks ResourceName ESXDetails -delete -keys
```

- 3 Encrypt the password of the vCenter Server user account.

For details on encrypting agent passwords, see the *Cluster Server Administrator's Guide*.

- 4 Specify the vCenter Server details.

```
hares -modify <VMwareDisks ResourceName> ESXDetails  
-add <vCenter IP address or hostname> <UserName>=<encrypted password>
```


Administering application monitoring

This chapter includes the following topics:

- [About the various interfaces available for performing application monitoring tasks](#)
- [Administering application monitoring using the Veritas High Availability tab](#)
- [Administering application monitoring settings](#)
- [Administering application availability using Veritas High Availability dashboard](#)

About the various interfaces available for performing application monitoring tasks

Veritas High Availability solution enables you to configure and administer application monitoring through vSphere Web Client and directly through a browser window.

Veritas High Availability provides you with the following interfaces to perform the application monitoring tasks through vSphere Web Client:

- Veritas High Availability tab
Use the Veritas High Availability tab to configure and administer application monitoring on a virtual machine.
- Veritas High Availability dashboard
Use the Veritas High Availability dashboard to administer application monitoring for multiple applications on multiple systems in a data center.

For information about how does the Veritas High Availability solution integrates with vSphere Web Client, see:

See [“How does the Veritas High Availability solution integrate with vSphere Web Client”](#) on page 10.

To perform the tasks directly through a browser window, use the following URL:

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

A prompt for the user account details will be displayed. You must enter the system user account details.

Administering application monitoring using the Veritas High Availability tab

Veritas High Availability provides you with an interface, the **Veritas High Availability** tab, to configure and control application monitoring. The **Veritas High Availability** tab is integrated with the VMware vSphere Web Client.

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. For more information on the later:

See [“Administering application availability using Veritas High Availability dashboard”](#) on page 49.

Use the **Veritas High Availability** tab to perform the following tasks:

- 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 resolve a held-up operation
- To modify application monitoring settings
- To view application dependency

- To view component dependency

To view the **Veritas High Availability** tab, launch the VMware vSphere Web Client, select a virtual machine from the inventory and then click the **Veritas High Availability** tab.

Understanding the Veritas High Availability tab work area

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

When you click a system in the inventory view of the VMware vSphere Client, the **Veritas High Availability** tab displays application information for the entire VCS cluster, not just the selected system.

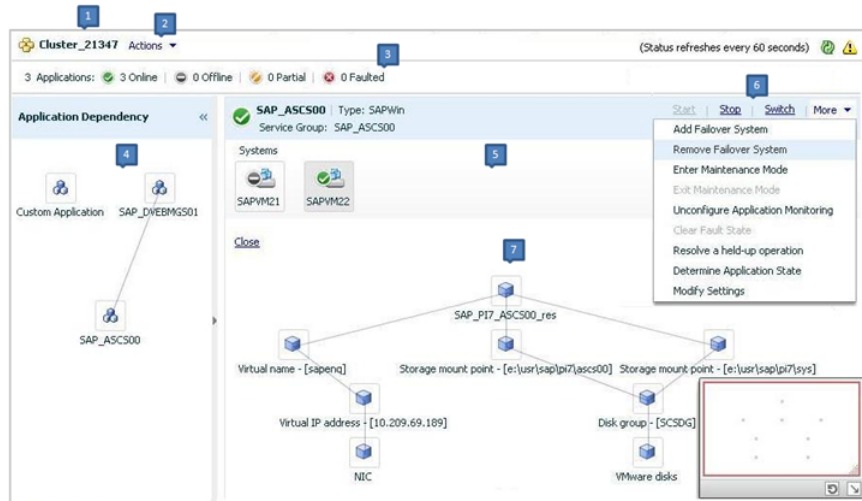
Note: If you do not configure any application for monitoring in the VCS cluster, then the **Veritas High Availability** tab displays only the following link: **Configure an application for high availability**.

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

The **Veritas High Availability** tab displays complex applications, like SAP Netweaver, in terms of multiple interdependent instances of that application. These interdependent instances represent component groups of the application. The component groups are also known as "service groups" in VCS terminology.

Each service group in turn includes several critical components of the application. The components are known as "resources" in VCS terminology.

The following figure displays two instances of SAP running in the **Veritas High Availability** tab:



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** tab graphic user interface (GUI) includes the following components:

- 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 virtual machine, 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 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 summary includes the total number of applications, and the state-wise breakdown of the applications in terms of the 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 former application appears at a lower level. 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 that are configured in the VCS cluster that is associated with the system you selected in the inventory view.
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 the 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 drop-down list of more tasks. The More drop-down 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 complex 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 or 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 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>

To configure or unconfigure application monitoring

Use the **Veritas High Availability** tab to configure or unconfigure an application for monitoring in a cluster under Cluster Server (VCS) control.

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

- 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** tab 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.
- 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 the applications that are 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.

To start or stop applications

Use the following options on the **Veritas High Availability** tab 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 predefined 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 or 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**.

Cluster Server 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 the systems that are part of a VCS cluster in a VMware environment. 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:

- Cluster Server (InfoScale Availability)7.3.1 is installed on the system.
- The system is not part of any other VCS cluster.
- The system has at least two network adapters.
- 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 virtual machine. 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.
Domain/Username	<p>Specify the user name with administrative privileges on the system.</p> <p>Specify the user name must be in the <i>domain.com\username</i> format.</p> <p>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.</p>
Password	Specify the password for the account you specified.
Use the specified user account on all systems	This option is checked by default. You cannot modify this setting.

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

- 4 Specify the user name and that VCS agents must use to perform domain operations such as Active Directory updates.
- 5 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	<p>Select the IP address to be used for cluster communication over the specified UDP port.</p>
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	<p>Displays the subnet mask to which the specified IP belongs.</p>

- 6 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.
Virtual name	Specifies a virtual name.
NIC	For each newly added system, specify the network adaptor that must host the specified virtual IP.

- 7 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** dialog box, specify the following details, and then click **Next**:

ESX host name or IP address	Specify the target ESX host name or IP address. The virtual machines can failover 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 that is 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.

- 8 On the **Configuration Summary** panel, review the VCS cluster configuration summary, and then click **Next** to proceed with the configuration.

- 9 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 tab.
 - Click **Next**.
- 10 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 *Cluster Server 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** tab 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** tab, 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 tab 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** tab displays the consolidated health information of all applications that are configured for monitoring in a VCS cluster. The tab automatically refreshes the application health information 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** tab, 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 delete all cluster configurations, (including networking and storage configurations), as well as application-monitoring configurations.

- On the Title bar of the Veritas High Availability tab, 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. For non-secure clusters, specify the user name and password credentials of a domain user with local administrative privileges on each VCS cluster node, and then click **OK**.

Administering application monitoring settings

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

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.

Note: The dashboard presents a unified view of monitored applications in a data center. It displays an application-centric view, not a product-centric view. If you have configured applications under more one Veritas High Availability product (VCS or ApplicationHA), then you cannot determine which application is under the control of which Veritas High Availability product. However, you can conclude that applications configured for failover are under VCS control. Applications configured for monitoring without a failover system may either be under VCS control or under ApplicationHA control.

To understand how to navigate across the dashboard:

See [“Understanding the dashboard work area”](#) on page 50.

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)

Apart from applications on systems running Cluster Server, the Veritas High Availability dashboard also displays applications running on Symantec ApplicationHA guests (versions 5.1 SP2 and later).

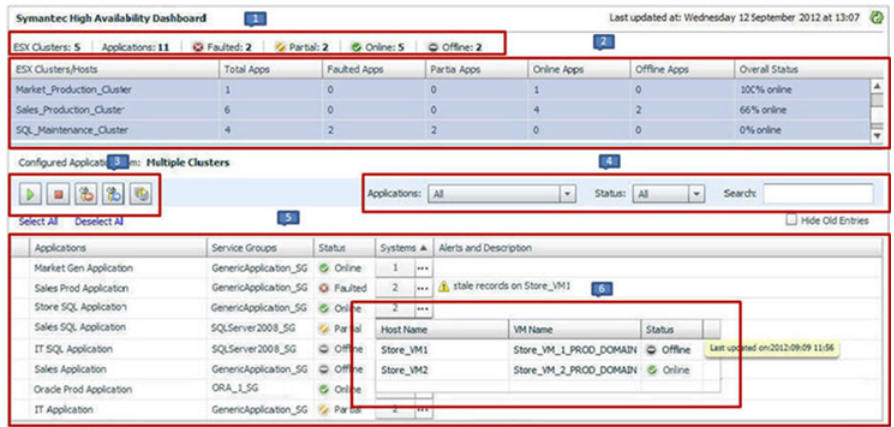
For more information on monitoring applications running on Symantec ApplicationHA, refer to Symantec ApplicationHA documentation.

Understanding the dashboard work area

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

Depending on whether you click a data center 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 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 (drop-down) |

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 data center 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. The equivalent term in Veritas High Availability terminology is “component group”.</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>
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	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 (drop-down) appears, listing the ESX host name of each configured system, the VM name (system name), and the status of the application on each system.

Column	Description
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>

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 data center or an ESX cluster. Various filters enable you 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 data center or ESX cluster. Click to select the name of the application that you want to filter.

Field	Description
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 54.

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 55.

- 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 54.

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 55.
- 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 drop-down 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 54.
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 55.
- 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 54.
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 55.
- 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 54.
 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 55.
- 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 failover applications to the faulted system, and application availability may be compromised. For more information, See [“To clear Fault state”](#) on page 46.

Troubleshooting

This appendix includes the following topics:

- [Troubleshooting application monitoring configuration issues](#)
- [Troubleshooting Veritas High Availability view issues](#)

Troubleshooting application monitoring configuration issues

This section lists common troubleshooting scenarios that you may encounter while or after configuring application monitoring.

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 the Veritas High Availability guest component (&ProdNameS_Enterprise;).

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 xprtld daemon is running on the system, restarting the xprtld 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 Veritas High Availability view issues

This section lists common troubleshooting scenarios that you may encounter when using the **Veritas High Availability** tab/dashboard.

Veritas High Availability tab not visible from a cluster node

If you click a system in the inventory view of the VMware vSphere Client GUI, then the **Veritas High Availability** tab displays the cluster view (consolidated cluster-level health information of the configured applications running on the selected system). In some multi-node cluster, the view is not visible from at least one of the cluster nodes.

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 node, and then retry.

Veritas High Availability tab does not display the application monitoring status

The **Veritas High Availability** tab in the vSphere Client console may either display a HTTP 404 Not Found error or may not show the application health status at all.

Verify the following conditions and then refresh the **Veritas High Availability** tab in the vSphere Client console:

- 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 (xprtld 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 xprtld
```
- 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 tab 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** tab may freeze
- The **Veritas High Availability** tab 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
```

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. At times, the view display may fail to load or fail to refresh itself after the default interval of 60 seconds. (2932028)

Workaround:

Restarting the xprtld service may resolve the issue.

To restart the xprtld service

- 1 Stop the xprtld service:

```
# /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:

```
# /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.