Veritas™ Cluster Server 6.0.4 Generic Application Agent Configuration Guide - Linux

October 2013

Veritas Cluster Server Generic Application Agent Configuration Guide

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

Product Version: 6.0.4

Document version: 6.0.4 Rev 1

Legal Notice

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

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations, whether delivered by Symantec 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.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

About Symantec Connect

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

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

Contents

Chapter 1

Introducing the Symantec High Availability solution for VMware

This chapter includes the following topics:

How the Symantec High Availability solution works in a VMware environment

The Symantec High Availability solution for VMware employs Veritas Cluster Server (VCS) and its agent framework to monitor the state of the 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 perfom the disk deatch and attach operations. However, instead of the ESX/ESXi hosts you can choose to communicate with the vCenter Server to perform these operations.

For more details refer to,

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

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.

For details on the VCS configuration concepts and clustering topologies, refer to the *Veritas 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 *VCS 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 deatch 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 VMareDisks agent from the target virtual machine sends a request to the vCenter Server to detach the disks from the failed virtual machine. However, since the host ESX has faulted, the request to detach the disks fails. The VMwareDisks agent from the target virtual machine now sends the disk attach request. The vCenter Server then processes this request and disks are attached to the target virtual machine. The application availability is thus not affected.

Limitation

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

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

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

Workaround

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

About the vCenter Server user account privileges

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

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

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

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

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

Typical VCS cluster configuration in a virtual environment

A typical VCS cluster configuration for generic applications, 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.

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

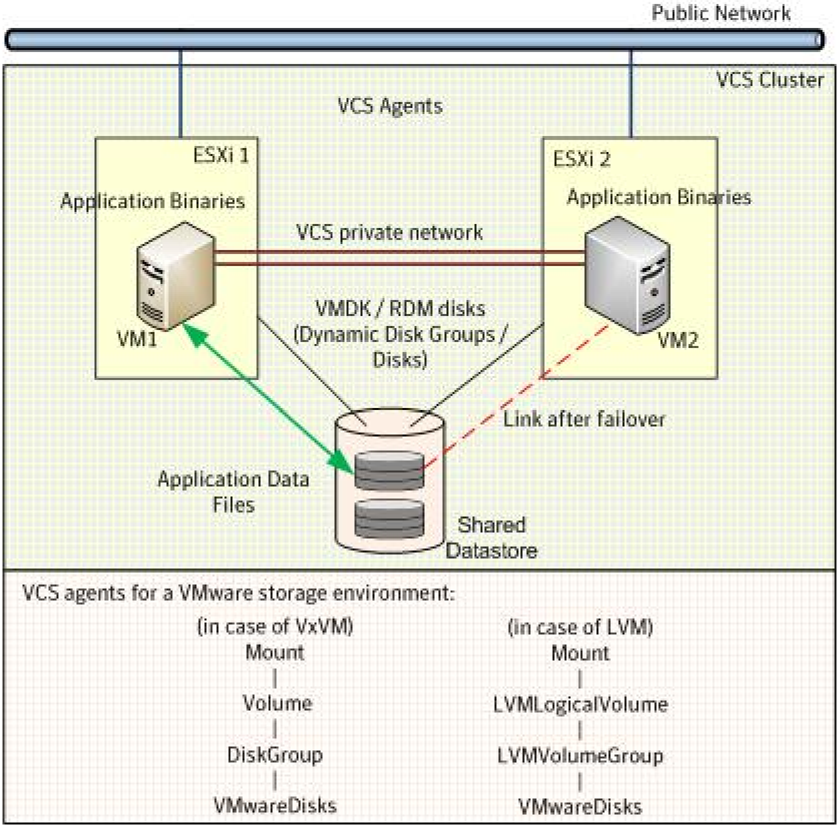

During a failover, the VCS storage agents move the VMware disks to the new system. The VCS network agents bring the network components online, and the application specific agents then start application services on the new system.

Chapter 2

About the Application agent

This chapter includes the following topics:

- About the Application agent for generic applications

About the Application agent for generic applications

The Application agent is a Veritas Cluster Server (VCS) bundled agent. This means that, when you install VCS on a virtual machine, the VCS Application agent is automatically installed on that machine.

The VCS Application agent is an agent that you can use to monitor high availability of generic applications. That is, applications that meet certain common criteria, such as the availability and accessibility of pre-defined start and stop programs. If an application-specific VCS agent exists, Symantec recommends the use of the specialized agent for monitoring availability. In other cases, Symantec recommends the use of the VCS Application agent.

The VCS Application agent brings generic applications online, takes them offline, and monitors their status. Use it to specify different executables for the online, offline, and monitor routines for different programs.

An application runs in the default context of root.

You can monitor the application in the following ways:

- Use the monitor program
- Specify a list of processes
- Specify a list of process ID files
- Any combination of the above.

The Application agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification.

For more information about the VCS Application agent, including descriptions of the agent functions and attributes, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

Configuring application monitoring for generic applications

This chapter includes the following topics:

Before configuring monitoring for generic applications

Ensure that you complete the following tasks before configuring monitoring for generic applications on a virtual machine:

- Install Veritas Cluster Server on the virtual machine that you need to monitor.
- Install and enable VMware Tools on the virtual machine where you want to monitor applications with VCS. Install a version that is compatible with the VMware ESX server.
- Install the VMware vSphere Client. The vSphere Client is used to configure application monitoring.

  You can also configure application monitoring directly from a browser window using the following URL:

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

  where, *VMNameorIP* is the name or IP address of the virtual machine on which you want to configure application monitoring.

- Install Symantec High Availability Console on a Windows system in your data center and register the Symantec High Availability plug-in with the vCenter server.
- Assign Configure Application Monitoring (Admin) privileges to the logged-on user on the virtual machine where you want to configure application monitoring.
- Install the application and the associated components that you wish to monitor on the virtual machine.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Veritas Cluster Server installer, wizards, and services.

  Refer to the *Symantec High Availability Solution Guide for VMware* for a list of ports and services used.
- If your application uses storage mount points, you must ensure that those mount points are already mounted on the virtual machine from which you are configuring the application for monitoring. All the required disks must be attached and all the storage components must be available.

  You must launch the Symantec High Availability Configuration Wizard from the virtual machine on which the application is running. The wizard discovers the disks that are attached and the storage that is currently available.
- 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 failover. 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.
- While creating a VCS cluster in a virtual environment, you must configure the cluster communication link over a public network in addition to private adapters. The link using the public adapter should be assigned as a low-priority link. This helps in case the private network adapters fail, leading to a condition where the systems are unable to connect to each other, consider that the other system has faulted, and then try to gain access to the disks, thereby leading to an application fault.
- You must not select teamed network adapters for cluster communication. If your configuration contains teamed network adapters, the wizard groups them as "NIC Group #N", where "N" is a number assigned to the teamed network adapters. 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:
  - SSO configuration failure.

  - The wizard may fail to discover the specified network adapters.
  - The wizard may fail to discover or validate the specified system name.

- 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 to the failover target system.
- Verify that the disks used by the application that you want to monitor are attached to non-shared controllers so that they can be deported from the system and imported to another system.
- If multiple types of SCSI controllers are attached to the virtual machines, then storage dependencies of the application cannot be determined and configured.
- The term ‘shared storage’ refers to the removable disks attached to the virtual machine. It does not refer to disks attached to the shared controllers of the virtual machine.
- If you want to configure the storage dependencies of the application through the wizard, the LVM volumes or VxVM volumes used by the application should not be mounted on more than one mount point path.
- The host name of the system must be resolvable through the DNS server or locally, using `/etc/hosts` file entries.
- To review the information about the functions, attributes, and resource type definition of the VCS Application agent, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

Configuring application monitoring by using the Symantec High Availability Configuration Wizard

To configure a generic application for monitoring by using the Symantec High Availability Configuration Wizard:

1. Launch the Symantec High Availability Configuration Wizard.
2. Review the information on the Welcome panel and click **Next**.
3. Select **Generic Application**, and then click **Next**.

4 On the Component Selection panel, enter a name for the component and click **Add Component**.

The component you added appears in the Component box.

5 Specify the following details to configure the component for monitoring:

- **Start program**: The complete path of the start program script.
- **Stop program**: The complete path of the stop program script.
- **Program to force-stop the application**: If you selected the **Specify application force-stop program** option, enter the complete path of the force-stop program script.

  Note: If you do not select the ‘force-stop’ option, VCS uses the stop program script that you specify for stopping the application.

- At least one or more of the following:
  - **Monitor program**: The complete path of the monitor program script.
  - **Application-related processes to monitor**: Names of the application processes that must be monitored.
  - **Application-generated PID files**: Path names of the process ID (PID) files of your application.
- **Enable intelligent monitoring for this application**: Select or clear this option to enable or disable intelligent monitoring for the application component. This option is selected by default. Symantec recommends that you enable intelligent monitoring of the application component.
- **User**: The user name. Ensure that you specify a valid user with adequate privileges on the virtual machine where you configure the application. Else, application monitoring may fail.

To remove a component from the Component box, use the Remove icon.

6 To specify more application components for monitoring, repeat step 4 to step 5. Else, click **Next**.

7 On the Storage Selection panel, select the appropriate mount points for the application instances that require storage, and click **Next**.

Note: The Storage Selection panel does not appear if shared storage is not attached to the virtual machine.

8 On the Define Start-Stop Order panel, to define the dependency between the components, select an application component from the **Parent Component** box and then select the components that it depends on from the **Depends on** box. When starting the application, the components are brought online in the defined order.

Note: The Define Start-Stop Order panel appears only when you have added more than one component for monitoring.

9 Click **Next**.

10 On the Configuration Inputs panel, use the Edit icon to specify the user name and password of the systems for the VCS cluster operations.

Cluster systems lists the systems included in the cluster configuration. **Application failover targets** lists the systems to which the application can fail over. Move the required systems to the Application failover targets list. Use the up and down arrow keys to define the priority order of the failover systems. The local system is selected by default for both, the cluster operations and as a failover target.

11 Skip this step if you do not want to add more systems to your cluster.

To add a system to the cluster, on the Configuration Inputs panel, click **Add System**. In the Add System dialog box, specify the following details of the system that you want to add to the VCS cluster and click **OK**:

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

The wizard validates the details, and the system then appears in the Cluster Systems list.

To remove a system from the cluster or from the Application failover targets list, use the Remove icon.

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

If you want to modify the security settings for the cluster, click **Advanced Settings**. In the Advanced settings dialog box, specify the following details and click **OK**.

Use single sign-on authentication	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 username and password and click **OK**.

Note: The **Advanced Settings** link is not visible if the cluster is already created.

13 Click **Next**.

14 On the Network Details panel, select the type of network protocol to configure the VCS cluster network links and then specify the adapters for network communication. By default, the links are configured over Ethernet.

Note: Symantec recommends that 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.

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

- **Use MAC address for cluster communication (LLT over Ethernet)** : Select the adapter for each network communication link. You must select a different network adapter for each communication link. This communication type 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)**: Select the type of IP protocol and then specify the required details for each communication link. This communication type configures the links over the routed network. Choose this mode if the failover target systems reside in the same or different subnets. The adapters that you select must have an IP address. Symantec recommends that the IP address assigned to these adapters should be in different subnets.

Select the IP protocol (IPv4 or IPv6) and then specify the following:

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

By default, one of the links is configured as a low-priority link on a public network interface. The second link is configured as a high-priority link. To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

15 Click **Next**.

16 Skip this step if you do not want to specify a virtual IP address for the application component. In this case, you must delete the Subnet mask details that the wizard automatically populates in the **Subnet Mask** field and click **Next**

On the Virtual Network Details panel, select the IP network (IPV4 or IPv6). The IPv4 protocol is selected by default.

Select the appropriate component and specify the following details for each failover system:

Virtual IP address	Specify a unique virtual IP address. You can specify only one virtual IP address for each component.
Subnet Mask (for IPv4)	Specify the subnet mask details.
Prefix (for IPv6)	Select the prefix details.
Network Adapter	Select the network adapter that will host the virtual IP.

If you want to add another virtual IP address, click **Add virtual IP address**.

17 Click **Next**.

18 If you selected mount points for your application in step 7, the Target ESX Details panel appears.

On the Target ESX Details panel, specify all the ESX hosts to which virtual machines can fail over. Each ESX host must be able to access the required shared datastores that contain visible disks.

To specify the ESX hosts, click **Add ESX Host** and in the Add ESX Host dialog box, specify the following details and click **OK**:

ESX hostname or IP address	Specify the target ESX hostname or IP address. The virtual machines can fail over to this ESX host during vMotion. All the additional ESX hosts should have access to the datastore on which the disks used by the application reside.
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 for the user account provided in the User name text box.

The wizard validates the user account and the storage details on the specified ESX hosts.

If you want to remove an ESX host, use the Remove icon.

19 Click **Next**.

20 On the Summary panel, review the VCS cluster configuration summary and then click **Next** to proceed with 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. Symantec recommends that you to validate the uniqueness of the assigned ID in the existing network. If the assigned ID is not unique or if you want to modify the cluster name or cluster ID, click **Edit**. In the Edit Cluster Details dialog box, modify the details as necessary and click **OK**.

21 On the Implementation panel, the wizard creates the VCS cluster, configures the application for monitoring, and creates cluster communication links.

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

If a configuration task fails, click **View Logs** to check the details of the failure. Rectify the cause of the failure and run the wizard again to configure application monitoring.

22 Click **Next** and then click **Finish** to complete the wizard workflow.

This completes the application monitoring configuration.

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

Sample configurations

This appendix includes the following topics:

- Sample configuration for an init process and generic application component

Sample configuration for an init process and generic application component

This section describes a sample procedure for using the Symantec High Availability Configuration Wizard to configure monitoring for the following across two virtual machines—vm1 and vm2:

- an init process, such as CUPS
- a generic application, MyApplication

As part of the configuration process, the wizard configures a 2-node cluster between the virtual machines vm1 and vm2 running on ESX hosts—ESXHost1 and ESXHost2, respectively.

Let us assume that the generic application (MyApplication) can be started, stopped, monitored, and forcibly stopped by using the following scripts, respectively:

- `start_MyComponent`
- `stop_MyComponent`
- `monitor_MyComponent`

  The `monitor_MyComponent` script is written to comply with the MonitorProgram attribute of generic applications. For more information, see the description of the Application agent attributes in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- `forcestop_MyComponent`

To configure application monitoring using the Symantec High Availability Configuration Wizard

1. Launch the VMware vSphere Client and connect to the VMware vCenter Server that hosts the virtual machine.
2. From the vSphere Client's Inventory view in the left pane, select ***vm1***.
3. From the vSphere Client's Management view in the right pane, click the **Symantec High Availability** tab.

   To configure single sign-on, go to 4. If you have already configured single sign-on during installation, go to 5.
4. On the Symantec High Availability View panel, specify the credentials of a user account that has administrative privileges on the virtual machine and click **Configure**.

   The Symantec High Availability console sets up a permanent authentication for the user account on that virtual machine.
5. On the Symantec High Availability View panel, click the **Configure application for high availability** link.

   This launches the Symantec High Availability Configuration Wizard.
6. Review the information on the Welcome panel and click **Next**.
7. Select **Generic Application**, and then click **Next**.
8. On the Component Selection panel, enter a name for the CUPS process, for example, cups_Program and click **Add Component**.

   The component you added (cups_Program) appears in the Component box.

9 Specify the following details to configure cups_Program for monitoring:

Start program	`/etc/init.d/cups start`
Stop program	`/etc/init.d/cups stop`
Program to force-stop the application	*Not specified* **Note:** If you do not select the 'force-stop' option, VCS uses the stop program script that you selected for stopping the application.
Monitor program	`/etc/init.d/cups status` **Note:** init processes such as CUPS, do not require special monitor scripts.VCS uses the status option of the init script for monitoring. However you can also use your own program scripts to monitor such processes.
Application-related processes to monitor	*cupsd -C /etc/cups/cupsd.conf*
Application-generated PID files	`/var/run/cupsd.pid`
Enable intelligent monitoring for this application	Selected by default to enable intelligent monitoring
User	*username*

Note: You must specify at least one or more of the following: Monitor program, application-related processes to monitor, application-generated PID files.

10 To configure MyApplication for monitoring, add a name for the My Application component, for example, MyComponent_Program, and click **Add Component**.

The component you added (MyComponent_Program) appears in the Component box.

11 Specify the following details to configure MyComponent_Program for monitoring:

Start program	*/home/user1/myapplication/bin/start_MyComponent*
Stop program	*/home/user1/myapplication/bin/stop_MyComponent*
Program to force-stop the application	Select **Specify application force-stop program** and enter `forcestop_MyComponent`.
Monitor program	*/home/user1/myapplication/bin/monitor_MyComponent*
Application-related processes to monitor	*Not specified*
Application-generated PID files	*Not specified*
Enable intelligent monitoring for this application	Selected by default to enable intelligent monitoring
User	*username*

12 Click **Next**.

13 If the MyApplication application requires storage, on the Storage Selection panel, select the appropriate mount points and click **Next**.

14 On the Define Start-Stop Order panel, you can define the relationship between the CUPS process and MyApplication.

To bring the CUPS process online first and then MyApplication, in the **Parent Component** list, select ***MyComponent_Program*** and in the **Depends on** box, select ***cups_Program***.

15 Click **Next**.

16 On the Configuration Inputs panel, the wizard lists vm1—the virtual machine from which you launched the wizard. The wizard also lists vm1 in the Application failover targets list. To add vm2 to the cluster, click **Add System**, and in the Add System dialog box, specify the following details for vm2:

System Name or IP address	*vm2*
User name	*username* Typically, this is the root user.
Password	*password*
Use the specified user account on all systems	Select to use the specified user account on all the cluster systems.

17. Click **Next**.
18. On the Network Details panel, select **Use MAC address for cluster communication (LLT over Ethernet)**. Specify two adapters for each virtual machine.
19. Click **Next**.
20. On the Virtual Network Details panel, select ***MyComponent_Program***, and then select **IPv4** and specify the following details for each failover system:

Virtual IP address	*IP address*
Subnet Mask	*Subnet mask*
Network Adapter	For vm1: *eth0* For vm2: *eth1*

21. Click **Next**.
22. On the Target ESX Details panel, click **Add ESX Host** to add details of ESXHost1. In the Add ESX Host dialog box, specify the following details and click **OK**:

ESX hostname or IP address	*ESXHost1*
User name	*esxhostuser1*
Password	*password*

The wizard validates the user account and the storage details on the specified ESX hosts.

23. Click **Add ESX Host** to add details of ESXHost2. On the Add ESX Host dialog box, specify the following details and click **OK**

ESX hostname or IP address	*ESXHost2*
User name	*esxhostuser2*
Password	*password*

24. Click **Next**.
25. On the Summary panel, review the VCS cluster configuration summary and then click **Next** to proceed with the configuration.

26 On the Implementation panel, the wizard creates the cluster, configures application monitoring, and creates cluster communication links. The wizard displays the status of each task. After all the tasks are complete, click **Next**.

27 Click **Finish** to complete the wizard workflow.

This completes the application monitoring configuration.

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

Appendix B

Sample scripts for generic application

This appendix includes the following topics:

Sample scripts to start, stop, and monitor a generic application

You can write your own scripts for the VCS Application agent to bring a generic application online, take the application offline, and monitor the status of the application.

You can also modify the following sample scripts and use them to start, stop, and monitor the application.

- Sample script to start a generic application:

  ```
  #!/bin/sh
  touch /tmp/sampleapp          # add any steps, if required
  exit 0
  ```

  You can modify the sample start script to suit the application requirements. If you save the start script with the name startsampleapp, then to bring the application online, the agent function runs the following command:

  ```
  su - user -c startsampleapp
  ```

- Sample script to stop a generic application:

```
#!/bin/sh
rm -f /tmp/sampleapp          # add any steps, if required
exit 0
```

You can modify the sample stop script to suit the application requirements. If you save the stop script with the name stopsampleapp, then to bring down the application, the agent function runs the following command:

`su - user -c stopsampleapp`

Note: The value of the return code for the start and stop scripts must be 0. No other return codes are supported.

- Sample script to monitor a generic application:

```
#!/bin/sh
APPLICATION_IS_ONLINE=110
APPLICATION_IS_OFFLINE=100
if [ -f /tmp/sampleapp ] ; then        # add any steps, if required
exit $APPLICATION_IS_ONLINE
else
exit $APPLICATION_IS_OFFLINE
fi
```

If you save the monitor script with the name monitorsampleapp, then to monitor the application, the agent function runs the following command:

`su - user -c monitorsampleapp`

About the monitor script exit codes

Custom monitor scripts use exit codes to let VCS know the status of the resource or process that is being monitored. The values that VCS expects as return values are:

- 1 or 100 - indicates that the resource is offline.
- 101 to 109 - indicates that the resource is online and has a confidence level of less than 100.
- 0 or 110 - indicates that the resource is online and has a confidence level of 100.

If the exit value returned is not one of the values listed above, then the status is considered unknown (typically a value of 99 is used).