About the Application agent

This section includes the following topics:

- About Generic Application
- About the Application agent for generic applications

About Generic Application

A generic application is an application which can be configured by defining:

- Start program
- Stop program
- (Optional) Force stop program
- One or more monitoring options like monitor program, list of processes, or process ID files

You can use the VCS Application agent to ensure high availability of a generic application.

If an application-specific VCS agent exists, Veritas recommends the use of the specialized agent for monitoring availability. In other cases, Veritas recommends the use of the VCS Application agent.

Cluster Server 8.0 Generic Application Agent Configuration Guide - AIX, Linux, Solaris - v8.0

About the Application agent for generic applications

The Application agent is a VCS bundled agent. This means that, when you install VCS on a physical machine, virtual machine, logical domain, or logical partition (LPAR), the VCS Application agent is automatically installed on that machine.

The Application agent for generic applications is supported across the following platforms:

- AIX
- Linux
- Solaris

The VCS Application agent can be used to perform online, offline and monitor operations. Use it to specify different executables for the online, offline, and monitor routines for the different applications.

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 *Cluster Server Bundled Agents Reference Guide*.

For more information about the IMF, refer to the *Cluster Server Administrator's Guide*.

Configuring application monitoring for generic applications

This section includes the following topics:

- About configuring application monitoring using the Veritas InfoScale Availability solution
- Before configuring monitoring for generic applications
- Launching the High Availability Configuration wizard
- Configuring application monitoring by using the Veritas High Availability Configuration wizard

About configuring application monitoring using the Veritas InfoScale Availability solution

Consider the following before you proceed:

- Apart from the Veritas High Availability Configuration wizard, you can also configure application monitoring using the VCS commands. For more information, refer to the *Cluster Server Administrator's Guide*.
- Veritas recommends that you first configure application monitoring using the wizard before using VCS commands to add additional components or modify the existing configuration.

Apart from configuring application availability, the wizard also sets up the other components required for successful application monitoring.

- In VMware environment, you must not suspend a system if an application is currently online on that machine. If you suspend a system, 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.

> **Note:** For details about deploying, configuring, and administering the Veritas InfoScale Availability solution in VMware environment, refer to theVeritas InfoScale 8.0 Solutions Guide for VMware.

Before configuring monitoring for generic applications

Ensure that you complete the following tasks before configuring monitoring for generic applications:

- Install the Cluster Server on the physical machine, virtual machine, logical domain, or LPAR on which you want to configure the application for monitoring.
- If you are going to launch the wizard from VOM, ensure that the cluster has been configured and running.
- Assign the following privileges to the logged-on user where you want to configure application monitoring:
    - When wizard is launched through vSphere client, assign Configure Application Monitoring (Admin) privileges.
    - When wizard is launched through VOM, the logged-on user group must be assigned the Admin role on the cluster or on the Availability perspective.

The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

- Install the application and the associated components that you want to monitor on the physical machine, virtual machine, logical domain, or LPAR.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Cluster Server installer, wizards, and services.

Verify that the following ports are not blocked by a firewall:

PHYSICAL ENVIRONMENT, LOGICAL DOMAIN, AND LPAR	5634, 14161, 14162, 14163, AND 14164
VMware environment	443, 5634, 14152, and 14153

Note: In the physical environment, logical domain, or LPAR, ensure that at least one of the following ports 14161, 14162, 14163, or 14164 is kept open.

- You must not select bonded interfaces for cluster communication. A bonded interface is a logical NIC, formed by grouping several physical NICs together. All NICs in a bond 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.

- 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 *Cluster Server Bundled Agents Reference Guide*.

You can download the latest documents from SORT: documents

- If your application uses storage mount points, you must ensure that those mount points are already mounted on the physical machine, virtual machine, logical domain, or LPAR 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 Veritas High Availability Configuration Wizard from the physical machine, virtual machine, logical domain, or LPAR on which the application is running. The wizard discovers the disks that are attached and the storage that is currently available.

Additional prerequisite for VOM

- The wizard option is available in VOM only after the cluster is configured and running, so configure the cluster through CPI installer or manually.

Additional prerequisites for VMware virtual environment

- 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/ESXi server.
- Install the VMware vSphere Client. You can configure application monitoring from the Veritas High Availability tab in the vSphere Client.

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

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

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

- Install Veritas High Availability Console on a Windows system in your data center and register the Veritas High Availability plug-in with the vCenter server.
- 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.
- Verify that the disks used by the application that you want to monitor are attached to non-shared controllers so that they can be detached from the system and attached to another system.
- 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 cannot 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 attach multiple types of SCSI controllers to the virtual machines because 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.
- 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.
- By default, the controller ID and port must remain the same on all cluster nodes. If you do not want the resource to have the same controller ID and port, you should localize the attribute for all cluster nodes. Localization allows all cluster nodes to have different controller IDs and port numbers. For more information about localizing an attribute, refer to the *Cluster Server Administrator's Guide.*

Launching the High Availability Configuration wizard

In a physical Linux environment, VMware virtual environment, logical domain, or LPAR, you can launch the High Availability Configuration wizard using:

- Veritas InfoScale Operations Manager Client: To launch the wizard from the Veritas InfoScale Operations Manager Client
- A browser window: To launch the wizard from a browser window

In addition, in a Linux environment, you can also launch the wizard using:

- haappwizard utility: To launch the wizard using the haappwizard utility

In addition, in a VMware virtual environment, you can also launch the wizard using:

- VMware vSphere Client: To launch the wizard from the VMware vSphere Client

To launch the wizard from the Veritas InfoScale Operations Manager Client

1. Log in to the Veritas InfoScale Operations Manager Management Server console.
2. In the Veritas InfoScale Operations Manager home page, click the **Availability** icon from the list of perspectives.
3. In the Data Center tree under the Manage pane, locate the cluster.
4. Navigate to the Systems object under the cluster.
5. Locate the system on which the application is running or application prerequisites are met.
6. Right-click on the system, and then click **Configure Application**.

To launch the wizard from a browser window

1. Open a browser window and enter the following URL:

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

system_name_or_IP is the system name or IP address of the system on which you want to configure application monitoring.

1. In the Authentication dialog box, enter the username and password of the user who has administrative privileges.
2. Depending on your setup, use one of the following options to launch the wizard:
    - If you have not configured a cluster, click the **Configure a VCS Cluster** link.

- If you have already configured a cluster, click **Actions > Configure Application for High Availability or theConfigure Application for High Availability** link.
- If you have already configured a cluster and configured an application for monitoring, click **Actions > Configure Application for High Availability**.

> **Note:** In the Storage Foundation and High Availability (SFHA) 6.2 release and later, thehaappwizardutility has been deprecated.

To launch the wizard using the haappwizard utility

The haappwizard utility allows you to launch the High Availability Configuration wizard. The utility is part of the product package and is installed in the /opt/VRTSvcs/bin directory.

- Enter the following command to launch the High Availability Configuration wizard:

```
happwizard [-hostname host_name] [-browser browser_name] [-help]
```

The following table describes the options of the happwizard utility:

Table: Options of the happwizard utility

- HOSTNAME	ALLOWS YOU TO SPECIFY THE HOST NAME OR IP ADDRESS OF THE SYSTEM FROM WHICH YOU WANT TO LAUNCH THE HIGH AVAILABILITY CONFIGURATION WIZARD. IF YOU DO NOT SPECIFY A HOST NAME OR IP ADDRESS, THE HIGH AVAILABILITY CONFIGURATION WIZARD IS LAUNCHED ON THE LOCAL HOST.
-browser	Allows you to specify the browser name. The supported browsers are Internet Explorer and Firefox. For example, enter `iexplore` for Internet Explorer and `firefox` for Firefox.
	Note: The value is case-insensitive.
-help	Displays the command usage.

To launch the wizard from the VMware vSphere Client

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 the virtual machine where you want to configure application monitoring.
3. Skip this step if you have already configured single sign-on during guest installation.

Select the High Availability tab and in the High Availability View page, specify the credentials of a user account that has administrative privileges on the virtual machine and click **Configure**.

The High Availability console sets up a permanent authentication for the user account on that virtual machine.

1. Depending on your setup, use one of the following options to launch the wizard:
    - If you have not configured a cluster, click the **Configure a VCS Cluster** link.
    - If you have already configured a cluster, click **Actions > Configure Application for High Availability or theConfigure Application for High Availability** link.
    - If you have already configured a cluster and configured an application for monitoring and to configure another application for monitoring, click **Actions > Configure Application for High Availability**.

Configuring application monitoring by using the Veritas High Availability Configuration wizard

Before you configure an application for monitoring, you must configure a cluster. Refer to the *VCS Installation Guide* for the steps to configure a cluster.

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

1. Launch the Veritas High Availability Configuration Wizard.
2. Review the information on the Welcome panel and click **Next**.
3. On the Application Selection panel, select **Generic Application and clickNext**.
4. On the Application Inputs panel, enter a name for the component and click **Add Component**.

The component you added appears in the Component box.

1. 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.
- **Force-stop program** : The complete path of the program script to forcefully stop 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. Veritas recommends that you enable intelligent monitoring of the application component.

For more information about intelligent monitoring, refer to the *Application agent* section in the *VCS Bundled Agents Reference Guide*.

- **User** : The user name. This user name is used to run all the agent scripts. Ensure that you specify a valid user with adequate privileges on the physical machine, virtual machine, logical domain, or LPAR where you configure the application. Else, application monitoring may fail.

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

1. To specify more application components for monitoring, repeat steps 4 and 5. Else, click **Next**.
2. 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 mounted on the physical machine, virtual machine, logical domain, or LPAR.

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

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.

Click **Next**.

1. If you have launched the wizard through VOM, go to step 12.

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

1. On the Configuration Inputs panel, 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.
    - Cluster systems lists the systems included in the cluster configuration.
    - Application failover targets lists the systems to which the application can fail over.
2. 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	Select to use the specified user account on all those
account on all systems	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.

1. Skip this step if you do not want to specify a virtual IP address for the application component.

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**.

Click **Next**.

1. This step is only applicable if you are configuring application monitoring on a virtual machine (VMware environment).

If you selected mount points for your application in step 7, the Storage HA Inputs panel appears.

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

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

ESX/ESXi hostname or IP address	Specify the target ESX/ESXi hostname or IP address.
	The virtual machines can fail over to this ESX/ESXi host during
	vMotion. All the additional ESX/ESXi 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/ESXi host. The user account must have administrator privileges on
	the specified ESX/ESXi 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/ESXi hosts.

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

Click **Next**.

1. On the Summary panel, review the configuration details information and click **Next** to proceed with the configuration.

If you want to rename the service group, click the **Rename** link and edit the service group name in the Rename service group name dialog box.

1. On the Implementation panel, the wizard performs the following tasks and displays the status of each task:
    - Configures the application for monitoring
    - In a VMware environment, creates cluster communication links

After all the tasks are complete, click **Next**.

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

1. Click **Next and then clickFinish** to complete the wizard workflow.

This completes the application monitoring configuration.

> **Note:** If the application is not running and if you want to start the application, clickStartto start the configured components on the system.

Sample configurations

This section 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 Veritas High Availability Configuration wizard.

To configure monitoring for the following across two virtual machines - Machine1 and Machine2:

- 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 machines Machine1 and Machine2 running on hosts - Host1 and Host2, 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 *Cluster Server Bundled Agents Reference Guide*.

- `forcestop_MyComponent`

To configure application monitoring using the Veritas High Availability Configuration wizard

1. Launch the Veritas 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 CUPS process, for example, cups_Program and click **Add Component**.

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

1. Specify the following details to configure the cups_Program for monitoring:

START PROGRAM	`/ETC/INIT.D/CUPS START`
Stop program	`/etc/init.d/cups stop`
Force-stop program	Not specified
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 . For example, root.

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

2. To configure MyApplication for monitoring, add a name for the MyApplication component, for example, MyComponent and click **Add Component**.

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

1. Specify the following details to configure MyComponent for monitoring:

START PROGRAM	`/MYAPPLICATION/BIN/START_MYCOMPONENT`
Stop program	`/myapplication/bin/stop_MyComponent`
Force-stop program	`/myapplication/bin/forcestop_MyComponent`
Monitor program	`/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 . This user name is used to run all the agent scripts. For example, root.

2. Click **Next**.
3. If the MyApplication application requires storage, on the Storage Selection panel, select the appropriate mount points and click **Next**.
4. 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* and in the Depends on box, select *cups_Program*.

1. Click **Next**.
2. On the Configuration Inputs panel, the wizard lists Machine1 - the machine from which you launched the wizard. The wizard also lists Machine1 in the Application failover targets list. To add Machine2 to the cluster, click **Add System** , and in the Add System dialog box, specify the following details for Machine2:

SYSTEM NAME OR IP ADDRESS	MACHINE2
User name	username
	Typically, this is the root user.

SYSTEM NAME OR IP ADDRESS	MACHINE2
Password	password
Use the specified user	Select to use the specified user account on all the
account on all systems	cluster systems.

3. Click **Next**.
4. On the Virtual Network Details panel, select ***MyComponent***, 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 Machine1: eth0
	For Machine2: eth1

5. Click **Next**.
6. On the Summary panel, review the VCS cluster configuration summary and then click **Next** to proceed with the configuration.
7. 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**.
8. 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 scripts for generic application

This section includes the following topics:

- Sample scripts to start, stop, and monitor a generic application
- About the monitor script exit codes

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).

Infrastructure service groups

This section includes the following topics:

- Infrastructure service groups

Infrastructure service groups

This section is applicable only if you are configuring application monitoring on a VMware virtual machine.

As part of configuring the application, the Veritas High Availability Configuration wizard:

- Configures application specific service groups and resources.
- Configures the VCS infrastructure service group (VCSInfraSG).

VCSInfraSG includes a resource called VCSNotifySinkRes. The type of this resource is Process. VCSNotifySinkRes configures and administers the notify_sink process on the guest. The notify_sink process sends the details about service groups and its attributes to the Veritas High Availability Console. This information is used for reporting purpose and is displayed on the Dashboard.

> **Note:** VCSInfraSG is an internal service group. You must not add or delete resources from this service group.

> The following are the VCSInfraSG notes:

- Before you configure the application for monitoring, ensure that SSO is configured between the Veritas High Availability Console and the guest. If SSO is not configured, VCSInfraSG fails to come online.
- If VCSInfraSG or VCSNotifySinkRes faults, ensure that SSO is configured between the Veritas High Availability Console and the guest. Clear the faults and bring the resource online again.
- VCSInfraSG or VCSNotifySinkRes must not be taken offline because it affects the information displayed on the Dashboard.