Veritas™ Cluster Server Bundled Agents Reference Guide

AIX

5.0 Maintenance Pack 3

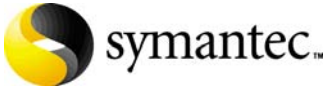

Veritas Cluster Server Bundled Agents Reference Guide

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: 5.0 MP3

Document version: 5.0MP3.0

Legal Notice

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
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 maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

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

www.symantec.com/techsupp

Contacting Technical Support

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

http://www.symantec.com/business/support/assistance_care.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/techsupp

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp

Customer Service is available to assist with 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 maintenance 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

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to clustering_docs@symantec.com.

Include the title and document version (located on the second page), and chapter and section titles of the text on which you

are reporting.

Maintenance agreement resources

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

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

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Chapter 5 Service and application agents

Chapter 6 Infrastructure and support agents

Chapter 7 Testing agents

Chapter 1

Introduction

Bundled agents are Veritas Cluster Server (VCS) processes that manage resources of predefined resource types according to commands received from the VCS engine, HAD. You install these agents when you install VCS.

A node has one agent per resource type that monitors all resources of that type. For example, a single IP agent manages all IP resources.

When the agent starts, it obtains the necessary configuration information from VCS. The agent then periodically monitors the resources, and updates VCS with the resource status.

Agents can:

- Bring resources online.
- Take resources offline.
- Monitor resources and report state changes.

For a more detailed overview of agents, see the VCS User's Guide.

Resources and their attributes

Resources are parts of a system and are known by their type, such as: a volume, a disk group, or an IP address. VCS includes a set of resource types. Different attributes define these resource types in the types.cf file. Each type has a corresponding agent that controls the resource.

The VCS configuration file, main.cf, contains the values for the resource attributes and has an include directive to the types.cf file.

An attribute's given value configures the resource to function in a specific way. By modifying the value of a resource attribute, you can change the way the VCS agent manages the resource. For example, the IP agent uses the Address attribute to determine the IP address to monitor.

Modifying agents and their resources

Use the Cluster Manager (Java Console), Veritas Cluster Server Management Console, or the command line to dynamically modify the configuration of the resources managed by an agent.

See the *Veritas Cluster Server User's Guide* for instructions on how to complete these tasks.

VCS enables you to edit the main.cf file directly. To implement these changes, make sure to restart VCS.

Attributes

Attributes contain data about the cluster, systems, service groups, resources, resource types, and the agent. An attribute has a definition and a value. You change attribute values to configure VCS resources. Attributes are either optional or required, although sometimes attributes that are optional in one configuration might be required in other configurations. Many optional attributes have predefined or default values, which you should change as required.

A variety of internal use only attributes also exist. Do not modify these attributes—modifying them can lead to significant problems for your clusters.

Attributes have type and dimension. Some attribute values can accept numbers, others can accept alphanumeric values or groups of alphanumeric values, while others are simple boolean on/off values.

Table 1-1 Attribute data types

Data Type	Description
string	Enclose strings, which are a sequence of characters, in double quotes ("). Optionally enclose strings in quotes when they begin with a letter, and contains only letters, numbers, dashes (-), and underscores (_). A string can contain double quotes, but the quotes must be immediately preceded by a backslash. In a string, represent a backslash with two slashes (//).
integer	Signed integer constants are a sequence of digits from 0 to 9. You can precede them with a dash. They are base 10. Integers cannot exceed the value of a 32-bit signed integer: 21471183247.
boolean	A boolean is an integer with the possible values of 0 (false) and 1 (true).

Table 1-2 Attribute dimensions

Dimension	Description
scalar	A scalar has only one value. This is the default dimension.
vector	A vector is an ordered list of values. Each value is indexed using a positive integer beginning with zero. A set of brackets ([]) denotes that the dimension is a vector. Find the specified brackets after the attribute name on the attribute definition in the types.cf file.
keylist	A keylist is an unordered list of unique strings.
association	An association is an unordered list of name-value pairs. An equal sign separates each pair. A set of braces ({}) denotes that an attribute is an association. Braces are specified after the attribute name on the attribute definition in the types.cf file, for example: str SnmpConsoles{}.

Chapter 2

Storage agents

This chapter contains:

About the storage agents

Use storage agents to Monitor shared storage.

DiskGroup agent

The DiskGroup agent brings online, takes offline, and monitors Veritas Volume Manager (VxVM) disk groups. This agent uses VxVM commands. You can use this agent to monitor or make disk groups highly available.

When the value of the StartVolumes and StopVolumes attribute is 1, the DiskGroup agent brings the volumes online and takes them offline during the import and deport operations of the disk group.

When you use a volume set, set StartVolumes and StopVolumes attributes of the DiskGroup resource that contains the volume set to 1. If a file system is created on the volume set, use a Mount resource to mount the volume set.

The agent protects data integrity by disabling failover when data is written to a volume in the disk group.

For important information on this agent, refer to:

"DiskGroup agent notes" on page 29

Dependencies

The DiskGroup resource does not necessarily depend on any other resource.

Figure 2-1 Sample service group for a DiskGroup resource

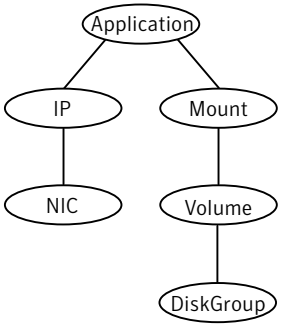

Agent functions

Online	Imports the disk group using the `vxdg` command.
Offline	Deports the disk group using the `vxdg` command.
Monitor	Determines if the disk group is online or offline using the `vxdg` command.

Clean

Terminates all ongoing resource actions and takes the resource offline—forcibly when necessary.

Info

The DiskGroup info agent function gets information from the Volume Manager and displays the type and free size for the DiskGroup resource.

Initiate the info agent function by setting the InfoInterval timing to a value greater than 0.

In the following example, the info agent function executes every 60 seconds:

```
# haconf -makerw
# hatype -modify DiskGroup InfoInterval 60
```

The command to retrieve information about the DiskType and FreeSize of the DiskGroup resource is:

```
# hares -value diskgroupres ResourceInfo
```

Output includes:

```
DiskType sliced
FreeSize 35354136
```

Action

Different action agent functions follow:

- license.vfd
  Checks for valid Veritas Volume manager license–if one is not found use the vxlicinst utility to install a valid license key.
- disk.vfd
  Checks if all disks in diskgroup are visible on host—if it fails, check if the path to disks exists from the host and check if LUN masking and zoning are set properly.
- udid.vfd
  Checks the UDIDs of disks on the cluster nodes—if it fails, ensure that the disks that are used for the disk group are the same on all cluster nodes.
- verifyplex.vfd
  Checks if the number of plexes on each site for the Campus Cluster setup are set properly—if it fails, check that the sites, disks, and plexes are set properly for a Campus Cluster setup.
- volinuse
  Checks if open volumes are in use or file systems on volumes that are mounted outside of VCS configuration.

See "High availability fire drill" on page 29.

State definitions

ONLINE	Indicates that the disk group is imported.
OFFLINE	Indicates that the disk group is not imported.
FAULTED	Indicates that the disk group has unexpectedly deported or become disabled.
UNKNOWN	Indicates that a problem exists either with the configuration or the ability to determine the status of the resource.

Attributes

Table 2-1 Required attributes

Required attribute	Description
DiskGroup	Name of the disk group that is configured with Veritas Volume Manager. Type and dimension: string-scalar

Table 2-2 Optional attributes

Optional attributes	Description
StartVolumes	If the value of this attribute is 1, the DiskGroup online function starts all volumes belonging to that disk group after importing the group. Type and dimension: string-scalar Default: 1
StopVolumes	If the value of this attribute is 1, the DiskGroup offline function stops all volumes belonging to that disk group before it deports the group. Type and dimension: string-scalar Default: 1

Table 2-2 Optional attributes

Optional attributes	Description
UmountVolumes	This attribute enables the DiskGroup resource to forcefully go offline even if open volumes are mounted outside of VCS control. When the value of this attribute is 1 and the disk group has open volumes, the following occurs: ■ The agent attempts to unmount the file systems on open volumes. If required, the agent attempts to kill all VCS managed and un-managed applications using the file systems on those open volumes. ■ The agent attempts to forcefully unmount the file systems to close the volumes. Type and dimension: integer-scalar Default: 0
TempUseFence	Do not use. For internal use only.
MonitorReservation	If the value of this attribute is 1, and SCSI-3 fencing is used, the agent monitors the SCSI reservation on the disk group. If the reservation is missing, the Monitor agent function takes the resource offline. Type and dimension: boolean-scalar Default: 0
NumThreads	The number of threads that are used within the agent process for managing resources. This number does not include the number of threads that are used for other internal purposes. Symantec recommends that you set the value of the NumThreads attribute to 1. Setting this attribute to a higher value may result in agent function timeouts due to serialization of underlying commands. Type and dimension: integer-scalar Default: 1

Table 2-2 Optional attributes

Optional attributes	Description
PanicSystemOnDGLoss	Determines whether to panic the node if the disk group becomes disabled. A loss of storage connectivity can cause the disk group to become disabled. If the value of this attribute is 1 and the disk group becomes disabled, the node panics. If the value of the attribute is 0 and the disk group becomes disabled, the following occurs: ■ If the cluster has I/O fencing enabled, the DiskGroup resource is marked FAULTED. This state results in the agent attempting to take the service group offline. As part of bringing the DiskGroup resource offline, the agent attempts to deport the disabled disk group. Even if disabled disk group fails to deport, the DiskGroup resource enters a FAULTED state. This state enables the failover of the service group that contains the resource. To fail back the DiskGroup resource, manually deport the disk group after restoring storage connectivity ■ If the cluster does not use I/O fencing, a message is logged and the resource is reported ONLINE. Type and dimension: boolean-scalar Default: 1

Resource type definition

```
type DiskGroup (
    static keylist SupportedActions = { "license.vfd", "disk.vfd",
    "udid.vfd", "verifyplex.vfd", checkudid, numdisks, campusplex,
    joindg, splitdg, getvxvminfo, volinuse }
    static int OnlineRetryLimit = 1
    static str ArgList[] = { DiskGroup, StartVolumes, StopVolumes,
    UmountVolumes, MonitorOnly, MonitorReservation, tempUseFence,
    PanicSystemOnDGLoss }
    str DiskGroup
    str StartVolumes = 1
    str StopVolumes = 1
    int UmountVolumes = 0
    static int NumThreads = 1
    boolean MonitorReservation = 0
    temp str tempUseFence = INVALID
    boolean PanicSystemOnDGLoss = 1
)
```

DiskGroup agent notes

The DiskGroup agent has the following notes:

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node. These discrepancies might prevent a service group from going online on a specific node.

For DiskGroup resources, the high availability fire drill checks for:

- The Veritas Volume Manager license
- Visibility from host for all disks in the disk group
- The same disks for the disk group on cluster nodes
- Equal number of plexes on all sites for the disk group in a campus cluster setup

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Setting the noautoimport flag for a disk group

VCS requires that the noautoimport flag of an imported disk group be explicitly set to true. This value enables VCS to control the importation and deportation of disk groups as needed when bringing disk groups online and taking them offline.

To check the status of the noautoimport flag for an imported disk group

- # **vxprint -l *disk_group* | grep noautoimport**
  If the output from this command is blank, the noautoimport flag is set to false and VCS lacks the necessary control.

For VxVM version 5.0 on AIX

The Monitor function changes the value of the VxVM noautoimport flag from off to on. It changes the value instead of taking the service group offline. This action allows VCS to maintain control of importing the disk group.

The following command changes the autoimport flag to false:

```
# vxdg -g disk_group set autoimport=no
```

For VxVM version 4.0

When you enable a disk group that is configured as a DiskGroup resource that does not have the noautoimport flag set to true, VCS forcibly deports the disk group. This forcible deportation may disrupt applications running on the disk group.

To explicitly set the noautoimport flag to true, deport the disk group and import it with the -t option as follows:

To deport the disk group, enter:

```
# vxdg deport disk_group
```

To import the disk group, specifying the noautoimport flag be set to true to ensure that the disk group is not automatically imported, enter:

```
# vxdg -t import disk_group
```

Configuring the Fiber Channel adapter

You must set FC adapter tunables appropriately to avoid excessive waits for monitor timeouts. One FS adapter tunable is FC error recovery policy.

Refer to the Fiber Channel adapter's configuration guide for further information.

Sample configurations

DiskGroup resource configuration

Example of a disk group resource in the Share Out mode.

```
DiskGroup dg1 (
    DiskGroup = testdg_1
)
```

DiskGroupSnap agent

Use the DiskGroupSnap agent to perform fire drills in a campus cluster. The DiskGroupSnap agent enables you to verify the configuration and data integrity in a Campus Cluster environment (with VxVM stretch mirroring).

For more information on fire drills, refer to the *Veritas Cluster Server User's Guide*.

For important information about this agent, refer to:

"DiskGroupSnap agent notes" on page 34

Dependencies

The DiskGroupSnap resource does not depend on any other resources. The service group that contains the DiskGroupSnap agent has an offline local dependency on the application's service group.

Figure 2-2 Sample service group for a DiskGroupSnap resource

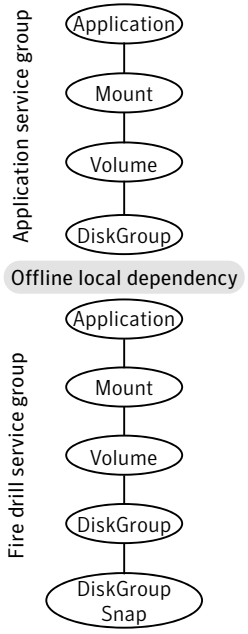

Agent functions

Online	Verifies that the application's disk group is in a valid campus cluster configuration. It detaches the site that the value of the FDSiteName attribute specifies. It then creates another disk group to be used for the fire drill on the detached site.
Offline	This re-attaches the site that the value of the FDSiteName attribute specifies back to the application's disk group.
Monitor	Monitors the DiskGroupSnap resource.
Clean	Takes the DiskGroupSnap resource offline.
Open	If the DiskGroupSnap resource has a parent resource that is not ONLINE, then it deletes the online lock file of the DiskGroupSnap resource. This marks the DiskGroupSnap resource as OFFLINE. In all other cases, the DiskGroupSnap resource performs no action.

State definitions

ONLINE	The DiskGroupSnap resource functions normally.
OFFLINE	The DiskGroupSnap resource is not running.
UNKNOWN	A configuration error exists.

Attributes

Table 2-3 Required attributes

Required attribute	Description
TargetResName	The name of the DiskGroup resource from the application's service group. Type-dimension: string-scalar Example: "dgres1"

Table 2-3 Required attributes

Required attribute	Description
FDSiteName	This is the site name that fire drill disks use. This name must be distinct for each site. You need to assign this local (per system) values as it maps to the SystemZone of the application service group. For more information about the SystemZone attribute, refer to the *Veritas Cluster Server User's Guide*. You can run the fire drill in the following two configurations: ■ Use a dedicated set of disks at the secondary that have been set aside for fire drill use. In this case, you must set the FDSiteName attribute to the VxVM site name given to this set of disks. This setup is commonly referred to as the Gold configuration. ■ Use the same disks that make up the mirror at the secondary site. In this case, you must set the FDSiteName attribute to the VxVM site name of the secondary site. This setup is commonly referred to as the Bronze configuration. Type and dimension: string-scalar Example: When the application service group has the following values for the SystemZones attribute: SystemZones = { n1 = 0, n2 = 0, n3 = 1, n4 = 1 } Where n1 (node 1) and n2 (node 2) comprise the first site and where the second site has n3 (node 3) and n4 (node 4). The FDSiteName definitions in the fire drill service group resemble the following: ■ "FDSiteName@n1=fdpri" ■ "FDSiteName@n2=fdpri" ■ "FDSiteName@n3=fdsec" ■ "FDSiteName@n4=fdsec" The fdpri and fdsec values are the site names of dedicated fire drill site disks at the primary and secondary sites respectively.

DiskGroupSnap agent notes

Configuration considerations

Keep the following recommendations in mind:

- Do not bring the DiskGroupSnap resource online in the SystemZone where the application service group is online.
- When you create the fire drill service group, in general use the same attribute values that you use in the application service group.
  However, the BlockDevice attribute of the Mount resource and the DiskGroup attribute of the DiskGroup resource change between the application's service group and the fire drill's service group. You must append an _fd to the original disk group name for the disk group name that the fire drill uses. For example, if dg01 is the disk group's name in the application service group, the attributes in the fire drill resemble those in Figure 2-3.
  Figure 2-3 shows the changes to resource values for the fire drill service group; note that the Volume resource is not included.

Figure 2-3 Sample resource values for a DiskGroupSnap resource

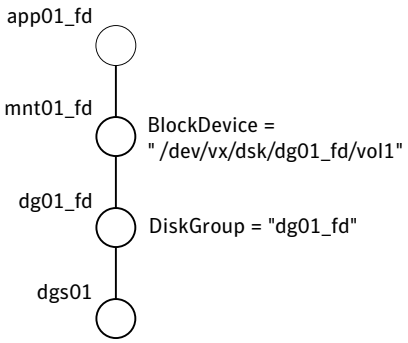

Agent limitations

The following limitations apply to the DiskGroupSnap agent:

- The online and offline operations of the DiskGroupSnap resource invokes VCS action entry points to run VxVM commands to detach/reattach the fire drill site. Since VxVM requires that these commands are run on the node where the disk group is imported, the disk group has to be imported on some node in the cluster before these operations.

- If you attempt to shut down the node where you brought the fire drill service group online, the node goes to a LEAVING state and the VCS engine attempts to take all the service groups offline on that node. At this point, the VCS engine rejects all action entry point requests. Therefore, during offline the DiskGroupSnap agent cannot invoke the action to reattach the fire drill site to the target diskgroup. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the fire drill site still remains detached from the diskgroup and you must manually reattach it.
- If you halt the node while the DiskGroupSnap resource's service group is still online, the VxVM site used for the fire drill remains detached after the node is brought up. You must manually reattach the fire drill site to the original diskgroup at the primary site.
- Before you shut down or stop VCS locally on the node where the fire drill service group is online, you must take the fire drill service group offline. Otherwise, after the node restarts you must manually reattach the fire drill site to the disk group that is imported at the primary site.

Resource type definition

```
type DiskGroupSnap (
    static int ActionTimeout = 120
    static int MonitorInterval = 300
    static int NumThreads = 1
    static str ArgList[] = { TargetResName, FDSiteName }
    str TargetResName
    str FDSiteName
)
```

Sample configurations

The following sample configure shows the fire drill's service group and its corresponding application service group. The fire drill's service group follows:

```
group dgfdsg (
    SystemList = { thoribm32 = 0, thoribm31 = 1 }
    SystemZones = { thoribm32 = 1, thoribm31 = 0 }
    )

    DiskGroup dgfdres (
        DiskGroup = newdg1_fd
        )

    DiskGroupSnap dgsres (
        TargetResName = dgres
```

```
        FDSiteName @thoribm32 = firedrill
        FDSiteName @thoribm31 = firedrill_31
        )

    Mount mntfdres1 (
        MountPoint = "/dgsfs1"
        BlockDevice = "/dev/vx/dsk/newdg1_fd/newvol1"
        FSType = vxfs
        FsckOpt = "-y"
        )

    Mount mntfdres2 (
        MountPoint = "/dgsfs2"
        BlockDevice = "/dev/vx/dsk/newdg1_fd/newvol2"
        FSType = vxfs
        FsckOpt = "-y"
        )

    Process procfdres1 (
        PathName = "/usr/bin/ksh"
        Arguments = "/scrib.sh /dgsfs1"
        )

    Process procfdres2 (
        PathName = "/usr/bin/ksh"
        Arguments = "/scrib.sh /dgsfs2"
        )

    requires group dgsg offline local
    dgfdres requires dgsres
    mntfdres1 requires dgfdres
    mntfdres2 requires dgfdres
    procfdres1 requires mntfdres1
    procfdres2 requires mntfdres2
```

The application's service group follows:

```
group dgsg (
    SystemList = { thoribm32 = 0, thoribm31 = 1 }
    SystemZones = { thoribm31 = 0, thoribm32 = 1 }
    )

    DiskGroup dgres (
        DiskGroup = newdg1
        )

    Mount mntres1 (
        MountPoint = "/dgsfs1"
        BlockDevice = "/dev/vx/dsk/newdg1/newvol1"
        FSType = vxfs
        FsckOpt = "-y"
        )
```

```
Mount mntres2 (
    MountPoint = "/dgsfs2"
    BlockDevice = "/dev/vx/dsk/newdg1/newvol2"
    FSType = vxfs
    FsckOpt = "-y"
    )

Process procres1 (
    PathName = "/usr/bin/ksh"
    Arguments = "/scrib.sh /dgsfs1"
    )

Process procres2 (
    PathName = "/usr/bin/ksh"
    Arguments = "/scrib.sh /dgsfs2"
    )

mntres1 requires dgres
mntres2 requires dgres
procres1 requires mntres1
procres2 requires mntres2
```

Volume agent

The Volume agent brings online, takes offline, and monitors a Veritas Volume Manager (VxVM) volume. You can use the agent to make a volume highly available or to monitor it.

Note: Do not use the Volume agent for volumes created for replication.

Dependencies

Volume resources depend on DiskGroup resources.

Figure 2-4 Sample service group for a Volume resource

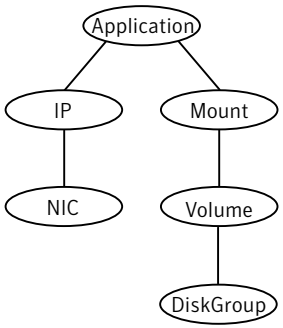

Agent functions

Online	Starts the volume using the `vxrecover` command.
Offline	Stops the volume using the `vxvol` command.
Monitor	Determines if the volume is online or offline by reading a block from the raw device interface to the volume.
Clean	Terminates all ongoing resource actions and takes the resource offline—forcibly when necessary.

State definitions

ONLINE	Indicates that the specified volume is started and that I/O is permitted.
OFFLINE	Indicates that the specified volume is not started and that I/O is not permitted.
FAULTED	Indicates the volume stops unexpectedly.
UNKNOWN	Indicates that the agent could not determine the state of the resource or that the resource attributes are invalid.

Attributes

Table 2-4 Required attributes

Required attribute	Description
DiskGroup	Name of the disk group that contains the volume. Type and dimension: string-scalar
Volume	Name of the volume from disk group specified in DiskGroup attribute. Type and dimension: string-scalar

Table 2-5 Optional attributes

Optional attributes	Description
NumThreads	Number of threads used within the agent process for managing resources. This number does not include threads used for other internal purposes. Symantec strongly recommends that you retain the default value of the NumThreads attribute of 1. Setting this attribute to a higher value may result in agent function timeouts due to serialization of underlying commands. Default: 1

Resource type definition

```
type Volume (
    static int NumThreads = 1
    static str ArgList[] = { Volume, DiskGroup }
    str Volume
    str DiskGroup
)
```

Sample configurations

Configuration

```
Volume v0 (
    Volume = vol0
    DiskGroup = testdg_1
)
```

LVMVG agent

The LVMVG agent activates, deactivates, and monitors a Logical Volume Manager (LVM) volume group. The LVMVG agent supports JFS or JFS2. It does not support VxFS. This agent ensures that the ODM is in sync with changes to the volume group. Specifically from the last time that the volume group was imported on the system.

For important information on this agent, refer to:

"LVMVG agent notes" on page 45

Dependencies

No dependencies exist for the LVMVG resource.

Figure 2-5 Sample service group for an LVMVG resource

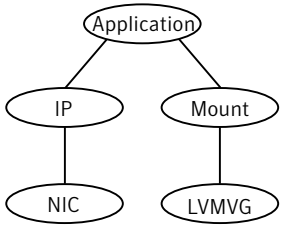

Agent functions

Online Activates the volume group. The Online agent function expects that the volume group is already imported on the system. If the volume group had been modified on a system where it was previously active, the online agent function detects the modification. It then syncs up the ODM on the system where you want to bring the volume group resource online.

Offline Deactivates the volume group.

Monitor Determines the volume group's state (activated or deactivated) and availability for read/write operations.

Clean Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

Action Different action agent functions follow:

- pv.vfd
  Checks if all the disks in the volume group are visible on a host. If it fails, check if the path to disks exists from the host and check if LUN masking and zoning are set properly.
- autoon.vfd
  Checks if the flag to automatically activate volume group on system restart is set to yes. If it fails, set the “auto on” flag of volume group to “no”.
- volinuse
  Checks if open volumes are in use or file systems on volumes that are mounted outside of VCS configuration.

State definitions

ONLINE	Indicates that the volume group is activated.
OFFLINE	Indicates that the volume group is deactivated.

Attributes

Table 2-6 Required attributes

Required attribute	Description
MajorNumber	Integer that represents the major number of the volume group. To ensure NFS functions properly, assign the same major number to the volume group on each system in the cluster. Type and dimension: integer-scalar
NumThreads	The number of threads that are used within the agent process for managing resources. This number does not include the threads that are used for other internal purposes. This resource type attribute is for internal use only. This value of this attribute must be set to 1. Type and dimension: integer-scalar Default: 1
VolumeGroup	Name of the volume group that is configured with LVM. Type and dimension: string-scalar Example: "testvg1"

Table 2-7 Optional attributes

Optional attribute	Description
GroupName	Attribute used to specify the volume's group. If set, the groups's name is applied to the volume group and all of its logical volumes. Type and dimension: string-scalar Default: system
ImportvgOpt	Attribute used to specify options for the importvg command. The default option, "n", indicates the volume group is not automatically activated when imported. Type and dimension: string-scalar Default: n
Mode	Attribute used to specify permissions for a volume group and its logical volumes. If set, these permissions are applied to the volume group and all of its logical volumes. Type and dimension: string-scalar Default: 640
OwnerName	Attribute used to specify the volume owner's name. If set, the owner's name is applied to the volume group and all of its logical volumes. Type and dimension: string-scalar Default: root

Table 2-7 Optional attributes

Optional attribute	Description
SyncODM	Integer that specifies whether or not the agent ensures that the ODM is in sync with any changes to the volume group. If the value of this attribute is 1, the agent ensures that the ODM is in sync with the changes to the volume group. In situations where the volume group was modified on another system in the cluster. The sync operation occurs on the system where the agent brings the volume group online. If the value of this attribute is 0, the changes to the volume group are independent of the ODM. See "SyncODM Attribute" on page 47. Type and dimension: integer-scalar Default: 1
VaryonvgOpt	Attribute used to specify options for the varyonvg command. Type and dimension: string-scalar

Resource type definition

```
type LVMVG (
    static keylist SupportedActions = { "pv.vfd", numdisks,
    "autoon.vfd", vo linuse }
    static int NumThreads = 1
    static str ArgList[] = { VolumeGroup, MajorNumber, OwnerName,
    GroupName,  Mode, ImportvgOpt, VaryonvgOpt, SyncODM }
    str VolumeGroup
    int MajorNumber
    str OwnerName
    str GroupName
    str Mode
    str ImportvgOpt = n
    str VaryonvgOpt
    int SyncODM = 1
)
```

LVMVG agent notes

The LVMVG agent for AIX has the following notes:

LVMVG support in a VIO server environment

The LVMVG agent supports volume groups created with virtual SCSI devices. AIX and VIOS must be at the following required levels:

- The AIX operating system level must be AIX 5.3 TL5 SP6 or later.
- The VIOS version must be VIOS 1.3 Fix Pack 8.1 or later.

Deactivation failure using the varyoffvg command on losing storage connectivity

In certain circumstances, the varyoffvg command does not deactivate all the volume groups on a node. This failure can prevent the failback of the LVMVG resource.

In situations where storage connectivity is lost, the LVMVG resources fails over. Failback for the LVMVG resource requires the deactivation of the volume groups on the node that lost its connectivity to storage. VCS uses the varyoffvg command to deactivate the volume groups. The LVMVG resource cannot fail back, however, when deactivation is unsuccessful.

When the volume group loses its storage connectivity, the clean function executes the `varyoffvg` command. Deactivation using the `varyoffvg` command can fail, however, if the volume group is busy. Criteria that can cause this failure can include:

- when the volume group has pending I/O operations, or
- when an application or upper-level resources in the resource dependency tree uses the volume group.

To overcome this deactivation failure, a post offline trigger has been added to issue the `varyoffvg` command. A side effect of the post offline trigger is that you must set the value of the OnlineRetryLimit attribute to 0.

After the restoration of storage connectivity, you must ensure that the volume groups are deactivated on the node. You can then clear the fault on the resources. If you find active volume groups, deactivate them using the varyoffvg command.

The LVMVG resource must be the bottom-most resource in the resource dependency tree in the service group. A resource under the LVMVG resource can potentially fail to go offline if the volume group's deactivation fails.

LVMVG Agent Supports JFS or JFS2

The LVMVG agent supports these file systems: JFS or JFS2. It does not support VxFS.

Volume group needs to be imported

The LVMVG agent relies on the ODM to find out the names of the disk devices that a volume group is created on. Unless a volume group is imported on the system, the ODM on that system does not contain any information about that volume group. Therefore, you must import the volume group on all the systems in the group's SystemList for the LVMVG agent to function properly.

For example, the volume groups (vg1 and vg2) must be imported on the specified systems (sysA and sysB).

See "LVMVG agent notes" on page 45.

Varyonvg options

By default, the agent checks the state of the disk devices underneath the volume group. If the disk device is in a defined state, the agent resets it to an available state. You can use the VaryonvgOpt attribute to change this default behavior.

You can tell the agent not to check for the state of the disk devices. Set the VaryonvgOpt attribute in the main.cf file to a value of "u". This option to the `varyonvg` command ensures that the disks underneath the volume group are not reserved when the volume group is activated.

Note: When you activate a volume group with the "u" option, ghost disks are not created. Therefore, you do not have to reset disks for these volume groups.

SyncODM Attribute

The LVMVG agent ensures that the ODM is in sync with any changes to the volume group since it was last imported on the system. This sync happens only if this attribute is set to 1. The agent maintains a time stamp file, /var/VRTSvcs/log/tmp/*volume_group_name*.ts, which records the time when the volume group was last imported on the system. When the agent initially brings a volume group online, the agent exports and reimports the group while initializing the time stamp file for that group. During the export and re-import processes, the agent preserves the ownership and mode information for the volume group and all its logical volumes.

The sync operation occurs when the time stamp value in the volume group's time stamp file is older than the time stamp value in the volume group's descriptor area. The timestamp value in the VGDA area of a volume group is updated after creating or deleting logical volumes, and adding or removing physical volumes.

Major Numbers

If a file system on a volume group is shared for NFS, make sure that the volume group is imported with the same major number. The volume group is imported on all of the nodes in the cluster.

To view a list of available major numbers on the system, enter the lvlstmajor command. For example:

```
# lvlstmajor
  49, 60 ...
```

To import volume group vg00 with major number 60, enter:

```
# importvg -V 60 -y vg00 hdisk3
```

To view the major number that is assigned to a volume group, use the `ls` command with the -l option. For example:

```
# ls -l /dev/vg00
crw-r-----   1 root     system    60,  0 Apr  2 16:05 /dev/vg00
```

Assign the same major number to the volume group on each system in the cluster. Specify this major number in the MajorNumber attribute of the LVMVG configuration.

Note: Do not specify the V option in the ImportvgOpt attribute string, the agent specifies this option.

Autoactivate Options

The "Concurrent Capable" options for the `importvg` and `mkvg` commands that are used with HACMP are not required for VCS. If an LVM volume group is placed under VCS control, the autoactivate options should be turned off. Do this using SMIT or through the command line.

From SMIT, set the following field values when creating or altering the volume group:

```
Activate volume group AUTOMATICALLY                 no
    at system restart?
Create VG Concurrent Capable?                       no
Auto-varyon in Concurrent Mode?                     no
```

From the command line, to view the current value for these fields, use the `lsattr` command.

For example:

```
# lsattr -El vg00
    vgserial_id  0001632f00004c00000000ee092b3bd8 N/A False
    auto_on      y                                N/A True
    conc_capable n                                N/A True
    conc_auto_on n                                N/A True
    timestamp    3ceff3390a8b1379                 N/A True
```

From the command line, to change the value for these fields, use the `chvg` command.

To change the value of `auto_on` to n:

1. Activate the volume group vg00 (if the volume group is not already activated):

   ```
   # varyonvg vg00
   ```

2. Run the `chvg` command:

   ```
   # chvg -a 'n' vg00
   ```

3 Verify the changes:

```
# lsattr -El vg00
    vgserial_id  0001632f00004c00000000ee092b3bd8 N/A False
    auto_on      n                                N/A True
    conc_capable n                                N/A True
    conc_auto_on n                                N/A True
    timestamp    3ceff3390a8b1379                 N/A True
```

LVMVG agent support for the Subsystem Device Driver (SDD)

The LVMVG agent supports the IBM Multipathing SDD version 1.4.0.0 and later. If disks are under SDD control, create a volume group with vpath devices. Refer to the SDD Documentation for configuration and migration of volume groups.

SDD support requires the /usr/sbin/lquerypr command, which provides a set of persistent reserve functions. The lquerypr command tool comes with the SDD installation package.

LVMVG agent support for the Hitachi's HiCommand Dynamic Link Manager (HDLM)

The LVMVG agent supports the Hitachi's HiCommand Dynamic Link Manager. For the details of the array and HDLM versions supported, refer to the HCL.

Note that if disks are under HDLM control, create a volume group with HDLM devices (dlmfdrv*n*). Refer to the HDLM documentation for configuration and migration of volume groups.

LVMVG agent support for the EMC PowerPath

The LVMVG agent supports the EMC PowerPath. For the details of the array and PowerPath versions supported, refer to the HCL.

Note that if disks are under PowerPath control, create a volume group with PowerPath devices (hdiskpower*n*). Refer to the EMC PowerPath documentation for configuration and migration of volume groups.

The hadevice utility

The LVMVG agent provides the hadevice utility. This utility checks the status of a disk device and resets a disk device to an available state. The utility then breaks any SCSI reservations on a disk device. Its syntax is:

```
hadevice -c | -r | -b -p device_name
```

The five possible states of a disk device are: AVAILABLE, DEFINED AND RESERVED, DEFINED AND UNRESERVED, PERSISTENT RESERVATION, and AVAILABLE AND OPEN.

To check the state of a disk device, enter:

```
# hadevice -c device_name
```

The following commands locate and remove ghost disks for a disk device and break any SCSI reservation on the disk device. When the -p flag follows the -b flag, it breaks any previous SCSI reservation on the device. It then obtains and retains a new reservation on the device. For SDD (vpath) disks, ghost disks are not created. Both the -b and -r flags remove any persistent reservation and clear all reservation key registration on the device. The -p flag (retain reservation) is not applicable for SDD disks.

To break any SCSI reservations on the disk device, enter:

```
# hadevice -b device_name
```

To break any SCSI reservations on the disk device, and obtain and retain a new reservation on the device, enter:

```
# hadevice -b -p device_name
```

To locate and remove ghost disks, reset a disk device that is in a DEFINED state and put it into an AVAILABLE state, enter:

```
# hadevice -r device_name
```

Removing a ghost disk from VxVM control

If VxVM 5.0 is installed, you may need to remove a ghost disk from VxVM control before using hadevice utility (except -r option).

If you check the ghost disk's status using the `hadevice -c hdisk#` command, you get an error. The error reads: `V-16-10011-10237 Error opening the device /dev/hdisk# (The file access permissions do not allow the specified action.)` Check if the ghost disk is under VxVM control. You can do this using the `vxdisk -eq list` command. If the disk is under VxVM control, remove it using the `vxdisk rm` *`vxvm_disk_name`*.

In this example, hdisk4 is a ghost disk.

```
sysA# vxdisk -eq list
Disk_0          auto      -          -         LVM        disk0
HDS9500-ALUA0_0 auto      -          -         error      hdisk4
HDS9500-ALUA0_1 auto      -          -         online     hdisk2
HDS9500-ALUA0_2 auto      -          -         online     hdisk3

sysA# vxdisk rm HDS9500-ALUA0_0
```

Sample configuration

```
system sysA

system sysB

system sysC

group  lvmgroup (
    SystemList = { sysA, sysB }
    AutoStartList = { sysA }

LVMVG lvmvg_vg1 (
    VolumeGroup = vg1
    MajorNumber = 50
    )

LVMVG lvmvg_vg2 (
    VolumeGroup = vg2
    MajorNumber = 51
    ImportvgOpt = "f"
    )
```

Mount agent

The Mount agent brings online, takes offline, and monitors a file system or an NFS client mount point. You can use the agent to make file systems or NFS mounted file systems highly available or to monitor them. This agent also supports high availability fire drills.

For important information about this agent, refer to:

"Mount agent notes" on page 59

Dependencies

No dependencies exist for the Mount resource.

Figure 2-6 Sample service group for a Mount resource

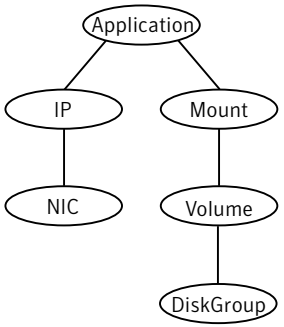

Agent functions

Online	Mounts a block device on the directory. If the mount process fails for non-NFS mounts, the agent attempts to run the `fsck` command on the device to remount the block device. If file system type is NFS, agent mounts the remote file system to a specified directory. The remote NFS file system is specified in the BlockDevice attribute.
Offline	Unmounts the mounted file system gracefully.
Monitor	Determines if the file system is mounted.
Clean	Unmounts the mounted file system forcefully.

Info

The Mount info agent function executes the command:

```
df -k mount_point
```

The output displays Mount resource information:

```
Size Used Avail Use%
```

To initiate the info agent function, set the InfoInterval timing to a value greater than 0. In this example, the info agent function executes every 60 seconds:

```
haconf -makerw
hatype -modify Mount InfoInterval 60
```

The command to retrieve information about the Mount resource is:

```
hares -value mountres ResourceInfo
```

Output includes:

```
Size 2097152
Used 139484
Available 1835332
Used% 8%
```

Action

- chgmntlock
  Invoke this action to reset the VxFS file system lock to a VCS-defined lock.
- mountpoint.vfd
  Checks if the specified mount point exists on the offline node. If it fails, it creates the mount point directory using `mkdir` command.
- mounted.vfd
  Checks if the mount point is already mounted on the offline node. If it fails, you need to unmount all the file systems from the specified mount point directory.
- vxfslic.vfd
  Checks for valid Veritas File System (VxFS) licenses. If it fails, you need to update the license for VxFS.
- mountentry.vfd
  Checks that the mount point is not listed in file system tables (e.g. /etc/filesystems).
  This action prevents the automatic mounting of the file system when the system reboots. If it fails, you need to remove mount point from file system tables.

State definitions

ONLINE	For the local file system, indicates that the block device is mounted on the specified mount point. For an NFS client, indicates that the NFS remote client is mounted on the specified mount directory.
OFFLINE	For the local file system, indicates that the block device is not mounted on the specified mount point. For an NFS client, indicates that the NFS remote client is not mounted on the specified mount directory.
FAULTED	For the local file system, indicates that the block device has unexpectedly unmounted. For the NFS client, indicates that the NFS remote client has unexpectedly unmounted.
UNKNOWN	Indicates that a problem exists either with the configuration or the ability to determine the status of the resource.

Attributes

Table 2-8 Required attributes

Required attribute	Description
BlockDevice	Block device for mount point. Type and dimension: string-scalar Example: "/dev/vx/dsk/myvcs_dg/myvol"
FsckOpt	Mandatory for non-NFS mounts. Use this attribute to specify options for the `fsck` command. You must correctly set this attribute for local mounts. If the mount process fails, the `fsck` command is executed with the specified options before it attempts to remount the block device. Its value must include either `-y`, `-n`, or `-p`. The -p option is only for jfs or jfs2 file systems on AIX. Refer to the fsck manual page for more information. For NFS mounts, the value of this attribute is not applicable and is ignored. Type and dimension: string-scalar Default: "-n" Example: `"-y"`
FSType	Type of file system. Supports jfs, jfs2, nfs, or vxfs. Type and dimension: string-scalar Example: "vxfs"
MountPoint	Directory for mount point Type and dimension: string-scalar Example: "/tmp/mnt"

Table 2-8 Required attributes

Required attribute	Description
VxFSMountLock	This attribute is only applicable to Veritas (VxFS) file systems. This attribute controls a file system locking feature to prevent accidental unmounts. This attribute can take three values: 0, 1, or 2. VxFSMountLock=0 The resource does not detect any changes to the lock when VCS reports that it is online after you set the value to zero. ■ If the mount point is initially locked with the mntlock="VCS", the monitor agent function unlocks it. ■ If the mount point is initially locked with a key that is not equal to "VCS", a message is logged once or the agent logs a message once. ■ If the mount point is initially not locked, no action is performed. VxFSMountLock=1 The resource does not detect changes to the lock when VCS reports it online after the value was set to one. VCS does not monitor the lock. ■ If the mount point is initially locked with the mntlock="VCS", no action is performed. ■ If the mount point is initially locked with a key that is not equal to "VCS", a message is logged once or the agent logs a message once. ■ If the mount point is initially not locked, the monitor agent function locks it with the mntlock="VCS". VxFSMountLock=2 When the value of the VxFSMountLock is 2, the file system is locked and the agent monitors any change to mntlock. ■ If the mount point is locked with the mntlock="VCS", no action is performed. ■ If the mount point is initially locked with a key that is not equal to "VCS", the monitor agent function logs a message whenever a change in mntlock is detected. ■ If the mount point is not locked, the agent locks it with the mntlock="VCS". Type and dimension: integer-scalar Default: 1

Table 2-9 Optional attributes

Optional attribute	Description
MountOpt	Options for the `mount` command. Refer to the `mount` manual page for more information. Do not set the VxFS mount option "mntlock=*key*". The agent uses this option only when bringing a Mount resource online. Type and dimension: string-scalar Example: "rw"
SnapUmount	If the value of this attribute is 1, this attribute automatically unmounts VxFS snapshots when the file system is unmounted. If the value of this attribute is 0, and snapshots are mounted, the resource cannot be brought offline. In this case, failover does not occur. Type and dimension: integer-scalar Default: 0
CkptUmount	If the value of this attribute is 1, this attribute automatically unmounts VxFS checkpoints when file system is unmounted. If the value of this attribute is 0, and checkpoints are mounted, then failover does not occur. Type and dimension: integer-scalar Default: 1
SecondLevelMonitor	This attribute is only applicable for an NFS client mount. It executes the `df -k` command for the NFS mounted file system and detects network outage. If the value of this attribute is 1, this attribute enables detailed monitoring of an NFS mounted file system. Type and dimension: boolean-scalar Default: 0

Table 2-9 Optional attributes

Optional attribute	Description
SecondLevelTimeout	This attribute is only applicable for an NFS client mount. This attribute is the timeout (in seconds) for the SecondLevelMonitor attribute that you try to request. The actual timeout value can be much smaller. This setting depends on how much time remains before it exceeds the MonitorTimeout interval. Type and dimension: integer-scalar Default: 30

Resource type definition

```
type Mount (
    static keylist SupportedActions = { "mountpoint.vfd",
    "mounted.vfd", "vxfslic.vfd", "mountentry.vfd", "chgmntlock" }
    static str ArgList[] = { MountPoint, BlockDevice, FSType,
    MountOpt, FsckOpt, SnapUmount, CkptUmount, SecondLevelMonitor,
    SecondLevelTimeout, VxFSMountLock }
    str MountPoint
    str BlockDevice
    str FSType
    str MountOpt
    str FsckOpt
    int SnapUmount = 0
    int CkptUmount = 1
    boolean SecondLevelMonitor = 0
    int SecondLevelTimeout = 30
    int VxFSMountLock = 1
)
```

Mount agent notes

The Mount agent has the following notes:

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node; discrepancies that might prevent a service group from going online on a specific node. For Mount resources, the high availability drill performs the following, it:

- Checks if the specified mount point directory exists
- Checks if the mount point directory is already used
- Checks for valid Veritas (VxFS) file system licenses
- Checks if the mount point exists in the /etc/filesystems file

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

VxFS file system lock

If the mount option in the mount table output has the option mntlock="*key*", then it is locked with the key "*key*". To verify if mount locking is in use and has the value of "*key*", run the `mount` command and review its output.

```
# mount
```

If the VxFS file system has mntlock="*key*" in its mount options, then unmounting the file system fails.

You can unlock the file system with the `fsadm` command and then unmount it. To unlock a locked mount, run the following command where "*key*" is the lock identifier and *mount_point_name* is the file system mount point.

```
# /opt/VRTS/bin/fsadm -o mntunlock="key" mount_point_name
```

To unmount a file system mounted with locking, run the `vxumount` command with the option mntunlock="*key*", for example:

```
# /opt/VRTS/bin/vxumount -o mntunlock="key" mount_point_name
```

Taking a group with the Mount resource offline can take several minutes if the file system is busy

When a file system has heavy I/O, the `umount` command can take several minutes to respond. However, the `umount` command temporarily deletes the mount point from mount command output while processing. Per IBM, this is the expected and supported behavior on AIX. The `umount` command's processing later puts the mount point back if the mount point is found busy. Meanwhile, the default OfflineTimeout value of the Mount agent can get exceeded, which in turn invokes the Clean agent function. The Clean function can find the mount point's entry absent from the mount command output and exit with success.

The unmounting, however, may not have happened yet. If unmounting did not occur, offlining resources below the Mount resource (for example the LVMVG or DiskGroup resources) can fail.

The Mount resource's Offline agent function then proceeds to unmount the mount point. After several attempts, the Clean scripts that clean the resources below the Mount resource succeed and the group goes offline.

See the *VCS User's Guide* for more information about the OfflineTimeout attribute.

Example 1

In this /etc/filesystems entry for a VxFS file system created on a VxVM volume, /mount_point is the mount point for the file system, /dev/vx/dsk/*Diskgroup_name*/*Volume_name* is the block device on which the file system is created, and `vxfs` is the file system type.

```
/etc/filesystems:
/mount_point:
    dev     = /dev/vx/dsk/Diskgroup_name/Volume_name
    vfs     = vxfs    mount    = false
    check   = false
```

Example 2

In this /etc/filesystems entry for a JFS file system created on an LVM logical volume, /mount_point2 is the mount point for the file system, /dev/*LVMlogical_volume* is the block device on which the file system is created, /dev/*LVMlogical_volumelog* is the log device for the file system automatically created by the `crfs` command, and jfs is the file system type.

```
/etc/filesystems:
/mount_point2:
           dev    = /dev/LVMlogical_volume
           vfs    = jfs
           log    = /dev/LVMlogical_volumelog
           mount  = false
           check  = false
```

Example 3

Use the `crfs` and `mkfs` commands to create file systems. VCS supports the following configurations for the Mount agent:

- LVM volume group with a JFS or JFS2 file system.
- VxVM volume with a VxFS file system.

Sample configurations

Configuration 1

In the following configuration, vg00 is a LVM volume group. The mount resource mnt requires the lvmvg_vg00 LVMVG resource.

```
LVMVG lvmvg_vg00 (
    VolumeGroup = vg00
    Disks = { "hdisk3" }
    Options = "u"
)

Mount mnt (
    MountPoint = "/lvm_testmnt"
    BlockDevice = "/dev/lv00"
    FSType = jfs
)
mnt requires vg00
```

Configuration 2

In the following configuration, vol0 is a volume in diskgroup testdg_1 created with VxVM. Mount resource m0 requires the dg1 diskgroup resource.

```
DiskGroup  dg1 (
    DiskGroup = testdg_1
)

Mount m0 (
    MountPoint = "/tmp/m0"
    BlockDevice = "/dev/vx/dsk/testdg_1/vol0"
    FSType = vxfs
)

m0 requires dg1
```

Configuration 3

In the following configuration, sysA is the remote NFS server and /home/xyz is the remote directory.

```
Mount mnt3 (
    MountPoint = "/tmp/m1"
    BlockDevice = "sysA:/home/xyz"
    FSType = nfs
)
```

ChapterX 3

Network agents

This chapter contains the following:

About the network agents

Use network agents to provide high availability for networking resources.

Agent comparisons

IP and NIC agents

The IP and NIC agents:

- Monitor a single NIC
- Support EtherChannel

IPMultiNIC and MultiNICA agents

The IPMultiNIC and MultiNICA agents:

- Monitor single or multiple NICs
- Check the backup NICs at fail over
- Use the original base IP address when failing over
- Provide slower failover compared to MultiNICB but can function with fewer IP addresses
- Have only one active NIC at a time

IPMultiNICB and MultiNICB agents

The IPMultiNICB and MultiNICB agents:

- Monitor single or multiple NICs
- Check the backup NICs as soon as it comes up
- Require a pre-assigned base IP address for each NIC
- Do not fail over the original base IP address
- Provide faster fail over compared to MultiNICA but require more IP addresses
- Have more than one active NIC at a time

802.1Q trunking

The IP/NIC, IPMultiNIC/MultiNICA, and IPMultiNICB/MultiNICB agents support 802.1Q trunking.

To use 802.1Q trunking, create 802.1Q trunked interfaces over a physical interface using SMIT. The physical interface is connected to a 802.1Q trunked port on the switch.

The NIC, MultiNICA, and MultiNICB agents can monitor these trunked interfaces. The IP, IPMultiNIC, and IPMultiNICB agents monitor the virtual IP addresses that are configured on these interfaces.

For example, create a 802.1Q interface called en6 over a physical interface called en0. Do not configure an IP address on en0. You connect en0 to a trunked port on the switch. The NIC and IP agents can then monitor en6 and the virtual IP address configured on en6.

IP agent

The IP agent manages the process of configuring a virtual IP address and its subnet mask on an interface. The virtual IP address must not be in use. You can use this agent when you want to monitor a single IP address on a single adapter.

The interface must be enabled with a physical (or administrative) base IP address before you can assign it a virtual IP address.

For the IP and NIC agents, VCS supports EtherChannel.

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node. These discrepancies might prevent a service group from going online on a specific node. For IP resources, the high availability fire drill checks for the existence of a route to the IP from the specified NIC.

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Dependencies

IP resources depend on NIC resources.

Figure 3-1 Sample service group for an IP resource

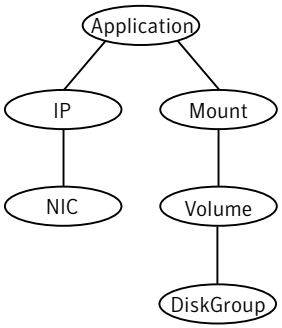

Agent functions

Online	Uses the `ifconfig` command to set the IP address as an alias on the interface.
Offline	Brings down the IP address that is specified in the Address attribute.
Monitor	Monitors the interface to test if the IP address that is associated with the interface is alive.
Clean	Brings down the IP address that is associated with the specified interface.

State definitions

ONLINE	Indicates that the device is up and the specified IP address is assigned to the device.
OFFLINE	Indicates that the device is down or the specified IP address is not assigned to the device.
UNKNOWN	Indicates that the agent could not determine the state of the resource or that the resource attributes are invalid.

Attributes

Table 3-1 Required attributes

Required attribute	Description
Address	A virtual IP address that is different from the base IP address, and that is associated with the interface. Note that the address you specify must not be the same as the configured physical IP address, but should be on the same network. Type and dimension: string-scalar Example: "192.203.47.61"
Device	The name of the NIC device that is associated with the IP address. Requires the device name without an alias. Type and dimension: string-scalar Example: "en0"
NetMask	The subnet mask that is associated with the IP address. Type and dimension: string-scalar Example: "255.255.255.0"

Table 3-2 Optional attributes

Optional attribute	Description
Options	Options for the `ifconfig` command. Type and dimension: string-scalar Example: "metric 4 mtu 1400"

Table 3-2 Optional attributes

Optional attribute	Description
RouteOptions	Specifies the routing options that are passed to the `route add` command when the agent configures an interface. The RouteOptions attribute value is generally formed like this: "*destination gateway metric*". For details about the `route` command, refer to the man page for your operating system. When the value of this string is null, the agent does not add routes. Type and dimension: string-scalar Example: "192.100.201.0 192.100.13.7" In this example, the agent executes the `"route add 192.100.201.0 192.100.13.7"` command when it configures an interface.

Resource type definition

```
type IP (
    static keylist SupportedActions = { "device.vfd", "route.vfd" }
    static str ArgList[] = { Device, Address, NetMask, Options }
    str Device
    str Address
    str NetMask
    str Options
)
```

Sample configurations

NetMask in decimal (base 10)

```
IP          IP_192_203_47_61 (
    Device = en0
    Address = "192.203.47.61"
    NetMask = "255.255.248.0"
    )
```

NetMask in hexadecimal (base 16)

```
IP          IP_192_203_47_61 (
    Device = en0
    Address = "192.203.47.61"
    NetMask = "0xfffff800"
    )
```

NIC agent

The NIC agent monitors the configured NIC. If a network link fails, or if a problem arises with the NIC, the resource is marked FAULTED. You can use the agent to make a single IP address on a single adapter highly available or to monitor it. This resource's Operation value is OnOnly.

For the NIC and IP agents, VCS supports EtherChannel.

The NIC listed in the Device attribute must have an administrative IP address. The administrative IP address is the default IP address that is assigned to the physical interface of a host on a network. This agent does not configure network routes or administrative IP addresses.

Before you use this agent:

- Verify that the NIC has the correct administrative IP address and subnet mask.
- Verify that the NIC does not have built-in failover support. If it does, disable it.

EtherChannel support

EtherChannel aggregates multiple network interfaces so that they appear as a single interface. For example, you can combine en0 and en1 into an EtherChannel and call the combined interface en2. You then use the NIC agent to monitor this en2 interface. You use the IP agent to configure and monitor an IP address on the en2 interface. Note that you use the en2 interface configured through EtherChannel for the Device attribute.

The IP and NIC agents support EtherChannel use with VCS. EtherChannel is responsible for providing local adapter swapping, which is outside of VCS control. EtherChannel Backup and active-active modes are supported.

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node. These discrepancies might prevent a service group from going online on a specific node. For NIC resources, the high availability fire drill checks for the existence of the NIC on the host.

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Dependencies

No child dependencies exist for this resource.

Figure 3-2 Sample service group for a NIC resource

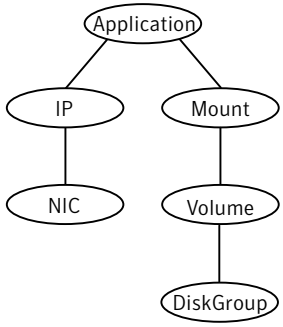

Agent functions

Monitor

Tests the network card and network link. Pings the network hosts or broadcast address of the interface to generate traffic on the network. Counts the number of packets passing through the device before and after the address is pinged. If the count decreases or remains the same, the resource is marked FAULTED.

If the NetworkHosts list is empty, or the ping test fails, the agent sends a ping to the device's broadcast address to generate network traffic. The agent checks for any response to the broadcast request. If there is no reply to the broadcast ping, the resource faults.

Note that for AIX, the systems do not respond to broadcast pings by default. Run the `no -o bcastping=1` command to enable response to broadcast pings.

State definitions

ONLINE	Indicates that the NIC resource is working.
FAULTED	Indicates that the NIC has failed.
UNKNOWN	Indicates the agent cannot determine the interface state. It may be due to an incorrect configuration.

Attributes

Table 3-3 Required attributes

Required attribute	Description
Device	Name of the NIC that you want to monitor. Use the `lsdev` command to check for all available network adapters. Type and dimension: string-scalar Example: "en0"
NetworkHosts	Required for virtual devices. See "NetworkHosts" on page 73.

Table 3-4 Optional attributes

Optional attribute	Description
NetworkHosts	List of hosts on the network that are pinged to determine if the network connection is alive. Enter the IP address of the host, instead of the host name, to prevent the monitor from timing out. DNS causes the ping to hang. If more than one network host is listed, the monitor returns ONLINE if at least one of the hosts is alive. If you do not specify network hosts, the monitor tests the NIC by sending pings to the broadcast address on the NIC. For a virtual device, you must configure the NetworkHosts attribute. Symantec recommends configuring more than one host to take care of the NetworkHost itself failing. Type and dimension: string-vector Example: { "166.96.15.22", "166.97.1.2" }
NetworkType	Type of network Type and Dimension: string-scalar Example: "ether"
PingOptimize	Determines whether to ping every monitor cycle. A value of 1 means that the agent pings either the network host or the broadcast address every monitor cycle. It pings each cycle to determine the state of the network interface. A value of 0 means that the agent uses the device statistics from the netstat output to determine the state of the interface. If no activity exists on the interface, the agent then pings the broadcast address to double-check the state of the network interface. Type and dimension: integer-scalar Default: 1

Resource type definition

```
type NIC (
    static keylist SupportedActions = { "device.vfd" }
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { Device, NetworkType, PingOptimize,
    NetworkHosts }
    static str Operations = None
    str Device
    str NetworkType
    int PingOptimize = 1
    str NetworkHosts[]
)
```

Sample configurations

Configuration without network hosts (using default ping mechanism)

```
NIC groupx_en0 (
    Device = en0
    PingOptimize = 1
    )
```

Configuration with network hosts

```
NIC groupx_en0 (
    Device = en0
    NetworkHosts = { "10.182.1.1", "10.182.1.2" }
    )
```

IPMultiNIC agent

The IPMultiNIC agent manages the virtual IP address that is configured as an alias on one interface of a MultiNICA resource. If the interface faults, the agent works with the MultiNICA resource to fail over to a backup NIC. If multiple service groups have IPMultiNICs associated with the same MultiNICA resource, only one group has the MultiNICA resource. The other groups have Proxy resources pointing to it. You can use this agent for IP addresses on multiple-adapter systems.

Dependencies

IPMultiNIC resources depend on MultiNICA resources.

Figure 3-3 Sample service group for an IPMultiNIC resource

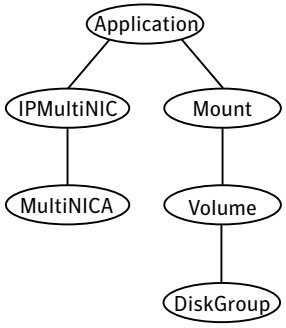

Agent functions

Online	Configures a virtual IP address on one interface of the MultiNICA resource.
Offline	Removes the virtual IP address from one interface of the MultiNICA resource.
Monitor	Checks if the virtual IP address is configured on one interface of the MultiNICA resource.

State definitions

ONLINE	Indicates that the specified IP address is assigned to the device.
OFFLINE	Indicates that the specified IP address is not assigned to the device.
UNKNOWN	Indicates that the agent can not determine the state of the resource. This state may be due to an incorrect configuration.

Attributes

Table 3-5 Required attributes

Required attribute	Description
Address	The virtual IP address that is assigned to the active NIC. Type and dimension: string-scalar Example: "10.128.10.14"
MultiNICAResName	Name of the associated MultiNICA resource that determines the active NIC. Type and dimension: string-scalar Example: "MultiNICA_grp1"
NetMask	Netmask for the virtual IP address. Type and dimension: string-scalar Example: "255.255.240.0"

Table 3-6 Optional attributes

Optional attribute	Description
Options	The `ifconfig` command options for the virtual IP address. Type and dimension: string-scalar Example: "mtu m"

Resource type definition

```
type IPMultiNIC (
        static str ArgList[] = { "MultiNICAResName:Device", Address,
NetMask, Options, "MultiNICAResName:Probed", MultiNICAResName }
    static int MonitorTimeout = 120
    str Address
    str NetMask
    str Options
    str MultiNICAResName
)
```

Sample configuration: IPMultiNIC and MultiNICA

Refer to the MultiNICA agent for more information.

```
group grp1 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )
    MultiNICA mnic (
        Device@sysa = { en0 = "10.128.8.42", en1 = "10.128.8.42" }
        Device@sysb = { en0 = "10.128.8.43", en1 = "10.128.8.43" }
        NetMask = "255.255.255.0"
        Gateway = "10.128.1.1"
        BroadcastAddr = "10.128.8.255"
        )

    IPMultiNIC ip1 (
        Address = "10.128.10.14"
        NetMask = "255.255.255.0"
        MultiNICAResName = mnic
        )

ip1 requires mnic
```

```
group grp2 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )

    IPMultiNIC ip2 (
        Address = "10.128.9.4"
        NetMask = "255.255.255.0"
        MultiNICAResName = mnic
        Options = "mtu m"
        )

Proxy proxy (
        TargetResName = mnic
        )

ip2 requires proxy
```

MultiNICA agent

The MultiNICA represents a set of network interfaces and provides failover capabilities between them. You can use the agent to make IP addresses on multiple-adapter systems highly available or to monitor them. Each interface in a MultiNICA resource has a base IP address. You can use one base IP address for all NICs, or you can specify a different IP address for use with each NIC. The MultiNICA agent configures one interface at a time. If it does not detect activity on the configured interface, it configures a new interface and migrates IP aliases to it.

If an interface is associated with a MultiNICA resource, do not associate it with any other MultiNICA, MultiNICB, or NIC resource. If the same set of interfaces must be a part of multiple service groups, configure a MultiNICA resource in one of the service groups. Configure the Proxy resources that point to the MultiNICA resource in the other service groups.

Dependencies

No dependencies exist for the MultiNICA resource.

Figure 3-4 Sample service group for a MultiNICA resource

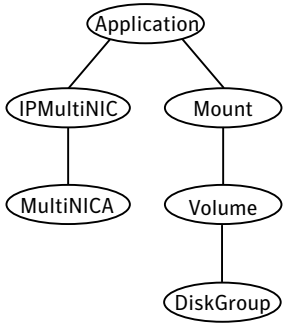

Agent function

Monitor — Checks the status of the active interface. If the agent detects a failure, it tries to migrate the IP addresses that are configured on that interface. If possible, it tries to migrate the addresses to the next available interface that is configured in the Device attribute.

Note: Systems do not respond to broadcast pings by default. You must run `"no -o bcastping=1"` to enable response to broadcast pings.

State definitions

ONLINE — Indicates that one or more of the network interfaces listed in the Device attribute of the resource is in working condition.

OFFLINE — Indicates that all of the network interfaces listed in the Device attribute failed.

UNKNOWN — Indicates that the agent cannot determine the state of the network interfaces that are specified in the Device attribute. This state may be due to incorrect configuration.

Attributes

Table 3-7 Required attributes

Required attribute	Description
BroadcastAddr	Broadcast address Type and dimension: string-scalar Example: "10.192.15.255"
Device	List of interfaces and their base IP addresses. Type and dimension: string-association Example: { en0 = "10.128.8.42", en1 = "10.128.8.42" }
Gateway	IP address for the default gateway. Type and dimension: string-scalar Example: "10.192.1.7"

Table 3-7 Required attributes

Required attribute	Description
NetMask	Netmask for the base IP address. Type and dimension: string-scalar

Table 3-8 Optional attributes

Optional attribute	Description
HandshakeInterval	Computes the maximum number of tries the agent makes either to: ■ ping a host (listed in the NetworkHosts attribute) when it fails over to a new NIC, or ■ ping the default broadcast address (depending on the attribute configured) when it fails over to a new NIC. To prevent spurious failovers, the agent must try to contact a host on the network several times before it marks a NIC as FAULTED. Increased values result in longer failover times, whether between the NICs or from system to system in the case of FAULTED NICs. Type and dimension: integer-scalar Default: 1
NetworkHosts	The list of hosts on the network that are pinged to determine if the network connection is alive. Enter the IP address of the host, instead of the host name, to prevent the monitor from timing out. DNS causes the ping to hang. If this attribute is unspecified, the monitor tests the NIC by pinging the broadcast address on the NIC. If more than one network host is listed, the monitor returns online if at least one of the hosts is alive. Type and dimension: string-vector Example: "128.93.2.1", "128.97.1.2"

Table 3-8 Optional attributes

Optional attribute	Description
Options	The `ifconfig` command options for the base IP address. Type and dimension: string-scalar Example: "metric 4 mtu 1400"
PingOptimize	Determines whether to ping every monitor cycle. A value of 1 means that the agent pings either the network host or the broadcast address every monitor cycle. It pings every cycle to determine the state of the network interface. A value of 0 means that the agent uses the device statistics from the netstat output to determine the state of the interface. If no activity exists on the interface, the agent then pings the broadcast address to double-check the state of the network interface. Type and dimension: integer-scalar Default: 1
RouteOptions	String to add a route when configuring an interface. The string contains the destination gateway metric. No routes are added if the value of this string is null. Type and dimension: string-scalar Example: "192.100.201.0 192.100.13.7"
FailoverInProgress	For internal use only.

Resource type definition

```
type MultiNICA (
    static int OfflineMonitorInterval = 60
    static int MonitorTimeout = 300
    static str ArgList[] = { Device, NetMask, Gateway,
    BroadcastAddr, Options, RouteOptions, PingOptimize,
    MonitorOnly, HandshakeInterval, NetworkHosts }
    static str Operations = None
    str Device{}
    str NetMask
```

```
    str Gateway
    str BroadcastAddr
    str Options
    str RouteOptions
    int PingOptimize = 1
    int HandshakeInterval = 1
    str NetworkHosts[]
    temp boolean FailoverInProgress = 0
)
```

MultiNICA notes

- If all NICs configured in the Device attribute are down, the MultiNICA agent faults the resource after a two-three minute interval. This delay occurs because the MultiNICA agent tests the failed NIC several times before it marks the resource OFFLINE. Failover logs record a detailed description of the events.
- The MultiNICA agent supports only one active NIC on one IP subnet; the agent does not work with multiple active NICs on the same subnet.
  - On AIX, for example, you have two active NICs, en0 (10.128.2.5) and en1 (10.128.2.8). You configure a third NIC, en2, as the backup NIC to en1. The agent does not fail over from en1 to en2 because some ping tests are redirected through en0 on the same subnet. The redirect makes the MultiNICA monitor return an online status.

EtherChannel support

EtherChannel aggregates multiple network interfaces so that they appear as a single interface. For example you can combine en0 and en1 into an EtherChannel and call the combined interface en2. You then use the MultiNICA agent to monitor this en2 interface. You use the IPMultiNIC agent to configure and monitor an IPMultiNIC address on the en2 interface. Note that you use the en2 interface configured through EtherChannel for the Device attribute.

The IPMultiNIC and MultiNICA bundled agents support EtherChannel use with VCS. EtherChannel is responsible for providing local adapter swapping, which is outside of VCS control. EtherChannel Backup and active-active modes are supported.

Sample configurations

MultiNICA and IPMultiNIC

In the following example, two systems, sysa and sysb, each have a pair of network interfaces, en0 and en1. In this example, the two interfaces, en0 and

en1, have the same base, or physical, IP address. Note the lines beginning Device@sysa and Device@sysb; the use of different physical addresses shows how to localize an attribute for a particular host.

The MultiNICA resource fails over the IP addresses to the backup NIC in the event of a failure of the active NIC. The resources ip1 and ip2, shown in the following example, have the Address attribute that contains the logical IP address. In the event of a NIC failure on sysa, the physical IP address and the two logical IP addresses fails over from en0 to en1.

However, if both the NICs on sysa are disconnected, the MultiNICA and IPMultiNIC resources work in tandem to fault the group on sysa. The entire group now fails over to sysb.

If you have more than one group using the MultiNICA resource, the other groups can use a Proxy resource. The Proxy resource points to the MultiNICA resource in the first group. The Proxy resource prevents redundant monitoring of the NICs on the same system. The IPMultiNIC resource is always made dependent on the MultiNICA resource.

See “IPMultiNIC agent” on page 75.

```
group grp1 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )
    MultiNICA mnic (
        Device@sysa = { en0 = "10.128.8.42", en1 = "10.128.8.42" }
        Device@sysb = { en0 = "10.128.8.43", en1 = "10.128.8.43" }
        NetMask = "255.255.255.0"
        Gateway = "10.128.1.1"
        BroadcastAddr = "10.128.25.255"
        Options = "mtu m"
        )

    IPMultiNIC ip1 (
        Address = "10.128.10.14"
        NetMask = "255.255.255.0"
        MultiNICAResName = mnic
        Options = "mtu m"
        )

ip1 requires mnic

group grp2 (
        SystemList = { sysa, sysb }
        AutoStartList = { sysa }
        )

        IPMultiNIC ip2 (
            Address = "10.128.9.4"
            NetMask = "255.255.255.0"
```

```
            MultiNICAResName = mnic
            Options = "mtu m"
            )
        Proxy proxy (
            TargetResName = mnic
            )

ip2 requires proxy
```

About the IPMultiNICB and MultiNICB agents

The IPMultiNICB and the MultiNICB agents can handle multiple NIC connections. Due to differences in the way that each platform handles its networking connections, these agents vary in design between platforms.

Checklist to ensure the proper operation of MultiNICB

For the MultiNICB agent to function properly, you must satisfy each item in the following list:

- Each interface must have a unique MAC address.
- At boot time, you must configure and connect all the interfaces that are under the MultiNICB resource and give them test IP addresses.
- All test IP addresses for the MultiNICB resource must belong to the same subnet as the virtual IP address.
- If you specify the NetworkHosts attribute, then that host must be on the same subnet as the other IP addresses for the MultiNICB resource.
- If any network host is meant to respond to a broadcast ping, run `no -o bcastping = 1` on the network host.
- You must use the AIX SMIT configuration tool to configure the test IP addresses and to make them persistent across reboots. If you do not use SMIT is to configure the IP addresses the agent may failover incorrectly.
- Ensure that media speed settings are the same for both the interface and the corresponding switch port. Symantec recommends setting the media speed to 100 Mbps full duplex.

IPMultiNICB agent

The IPMultiNICB agent works with the MultiNICB agent. The agent configures and manages virtual IP addresses (IP aliases) on an active network device that the MultiNICB resource specifies. When the MultiNICB agent reports a particular interface as failed, the IPMultiNICB agent moves the IP address to the next active interface. You can use this agent for IP addresses on multiple-adapter systems.

If multiple service groups have IPMultiNICB resources associated with the same MultiNICB resource, only one group should have a MultiNICB resource. The other groups should have a proxy resource pointing to the MultiNICB resource.

For the MultiNICB and IPMultiNICB agents, VCS supports EtherChannel.

Dependencies

IPMultiNICB resources depend on MultiNICB resources.

Figure 3-5 Sample service group for an IPMultiNICB resource

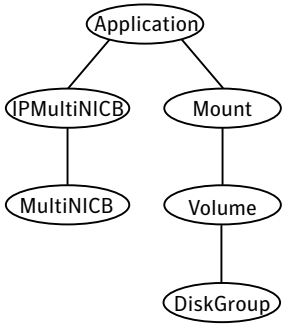

Requirements for IPMultiNICB

The following conditions must exist for the IPMultiNICB agent to function correctly:

- The MultiNICB agent must be running to inform the IPMultiNICB agent of the available interfaces.
- Only one IPMultiNICB agent can control each logical IP address.

Minimal configuration

The minimal configuration for this agent consists of:

- the failover IP address
- the subnet mask
- and the name of the MultiNICB resource that it depends on

See "Sample configurations" on page 91.

The haipswitch utility

You can use the haipswitch utility to switch IP addresses between MultiNICB interfaces on the same system. Running the utility with the `-h` flag gives an example of usage.

Agent functions

Online	Finds a working interface with the appropriate interface alias or interface name, and configures the logical IP address on it.
Offline	Removes the logical IP address.
Clean	Removes the logical IP address.
Monitor	If the logical IP address is not configured as an alias on one of the working interfaces under a corresponding MultiNICB resource, monitor returns OFFLINE. If the current interface fails, the agent fails over the logical IP address. It fails over the logical IP address to the next available working interface that is within the MultiNICB resource on the same node. If no working interfaces are available then monitor returns OFFLINE.
Open	Data structures necessary for monitoring the network interfaces are created.
Close	Data structures that the monitor agent function uses are freed.
Attr_Changed	Updates the data structures that are used for monitoring the NICs.

State definitions

ONLINE	Indicates that an IP address on one of the working network interfaces of the resource is up. The IP address is specified in the Address attribute. The resource is specified in the MultiNICBResName attribute.
OFFLINE	Indicates that an IP address on one of the working network interfaces of the resource is not up. The IP address is specified in the Address attribute. The resource is specified in the MultiNICBResName attribute.
UNKNOWN	Indicates that the agent cannot determine the status of the virtual IP address that is specified in the Address attribute.
FAULTED	The IP address could not be brought online, usually due to all NICs in the MultiNICB resource faulting.

Attributes

Table 3-9 Required attributes

Required attribute	Description
Address	Defines the dotted decimal failover IP address. This IP address must be different than the base or test IP addresses in the MultiNICB resource. The IPMultiNICB agent automatically assigns the failover IP address. Do not configure this IP address before the IPMultiNICB agent goes online. If the IP address is already configured, the agent returns an error. Type and dimension: string-scalar Example: "10.118.10.15"
MultiNICBResName	Contains the name of the MultiNICB resource that the IPMultiNICB resource depends on. Type and dimension: string-scalar Example: "MultiNICA_grp1"
NetMask	The netmask that is associated with the logical IP address. Type and dimension: string-scalar Example: "255.255.255.0"

Table 3-10 Optional attributes

Required attribute	Description
RouteOptions	Specifies the routing options that are passed to the `route add` command when the agent configures an interface. The RouteOptions attribute value is generally formed like this: "*destination gateway metric*". For details about the `route` command, refer to the man page for your operating system. When the value of this string is null, the agent does not add routes. Type and dimension: string-scalar Example: "192.100.201.0 192.100.13.7" In this example, the agent executes the `"route add 192.100.201.0 192.100.13.7"` command when it configures an interface.

Resource type definition

```
type IPMultiNICB (
    static int MonitorTimeout = 120
    static int OfflineMonitorInterval = 60
    static int MonitorInterval = 10
    static str ArgList[] = { Address, NetMask, MultiNICBResName,
    "MultiNICBResName:Probed"}
    str Address
    str NetMask
    str MultiNICBResName
)
```

Sample configurations

IPMultiNICB and MultiNICB

```
group grp1 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )

    MultiNICB MNICB_grp1 (
        Device@sysa = { en0 = "10.128.8.42", en1 = "10.128.8.43" }
        Device@sysb = { en0 = "10.128.8.44", en1 = "10.128.8.45" }
```

```
        NetworkHosts = "10.128.8.10"
        )

    IPMultiNICB ip1 (
        Address = "10.128.10.14"
        Netmask = "255.255.255.0"
        MultiNICBResName = MNICB_grp1
        )
    ip1 requires MNICB_grp1

group grp2 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )
    IPMultiNICB ip2 (
        Address = "10.128.10.15"
        Netmask = "255.255.255.0"
        MultiNICBResName = MNICB_grp1
        )
    Proxy MNICB_proxy (
        TargetResName = MNICB_grp1
        )
    ip2 requires MNICB_proxy
```

Other sample configurations for IPMultiNICB and MultiNICB

Refer to the sample configurations in the MultiNICB agent.

MultiNICB agent

The MultiNICB works with the IPMultiNICB agent. Allows IP addresses to fail over to multiple NICs on the same system before VCS tries to fail over to another system. You can use the agent to make IP addresses on multiple-adapter systems highly available or to monitor them.

When you use the MultiNICB agent, you must configure the NICs before putting them under the agent's control. You must configure all the NICs in a single MultiNICB resource with the IP addresses that are in the same subnet.

You need to set the MONITOR flag for each NIC that the agent controls. Use the `ifconfig` command to set the flag. For example:

```
# ifconfig en0 monitor
```

For the MultiNICB and IPMultiNICB agents, VCS supports EtherChannel.

EtherChannel support

EtherChannel aggregates multiple network interfaces so that they appear as a single interface. For example you can combine en0 and en1 into an EtherChannel and call the combined interface en2. You then use the MultiNICB agent to monitor this en2 interface. You use the IPMultiNICB agent to configure and monitor an IPMultiNICB address on the en2 interface. Note that you use the en2 interface configured through EtherChannel for the Device attribute.

The IPMultiNICB and MultiNICB bundled agents support EtherChannel use with VCS. EtherChannel is responsible for providing local adapter swapping, which is outside of VCS control. EtherChannel Backup and active-active modes are supported.

The haping utility

Use the haping utility (/opt/VRTSvcs/bin/ MultiNICB/haping) to test each NIC before you configure the MultiNICB resource. This utility takes the NIC interface as an argument. You can use the utility to perform a link test, a broadcast ping, or to ping a specific remote host. Symantec recommends that the administrator perform a test ping with the remote host before adding it to the NetworkHosts parameter. Some examples of the command syntax are as follows:

Link test only on interface en0:

```
haping -l en0
```

Ping a remote host 10.10.10.10 from interface en0:

```
haping -g 10.10.10.10 en0
```

Dependencies

No dependencies exist for the MultiNICB resource.

Figure 3-6 Sample service group for a MultiNICB resource

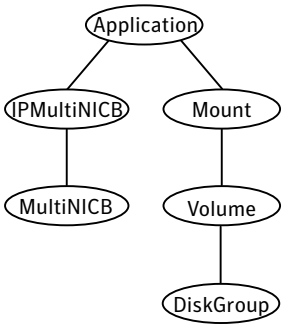

Agent functions

Open	Allocates an internal structure to store information about the resource.
Close	Frees the internal structure that is used to store information about the resource.
Monitor	Checks the status of each physical interface. Writes the status information to the export information file for IPMultiNICB resources to read it.

State definitions

ONLINE	Indicates that one or more of the network interfaces listed in the Device attribute of the resource is in working condition.
UNKNOWN	Indicates that the MultiNICB resource is not configured correctly.
FAULTED	Indicates that all of the network interfaces listed in the Device attribute failed.

Attributes

Table 3-11 Required attributes

Required attribute	Description
Device	Lists the interfaces that you want the agent to monitor. A unique test IP address must be assigned to each interface. You must use the AIX SMIT configuration tool to configure the test IP addresses and to make them persistent across reboots. **Note:** You also must manually configure the default IP route on each NIC in the MultiNICB resource. Type and dimension: string-association Example: { en1= "10.182.9.34", "en2=10.182.10.34" }
Gateway	IP address for the default gateway on the local network. Type and dimension: string-scalar Example: "136.22.1.1"

Table 3-12 Optional attributes

Optional attribute	Description
LinkTestRatio	Controls the frequency of the ping test in relation to the link test. The ping test may be run at a lesser frequency to reduce network traffic. If this attribute is set to 1, packets are sent during every monitor cycle. If this attribute is set to 0, packets are never sent during a monitor cycle. Symantec does not recommend setting the value to zero. The agent determines link status without transmitting any ping packets. For other values greater than 1, packets are sent at a lower frequency. For example, if LinkTestRatio=2, then ping packets are sent out during every other monitor cycle. In other words, packets are sent out half as often than if LinkTestRatio were equal to one. Type and dimension: integer-scalar Default: 1
NetworkHosts	The NetworkHosts attribute is a list of hosts on the local network that are pinged to determine if the network connection is available. These must be IP addresses, and not host names. If you do not specify this attribute, the agent monitors the NIC by pinging the broadcast address on the NIC. If you specify one or more network hosts, and at least one host responds to a ping, the agent reports the MultiNICB resource online. The IP addresses for the NetworkHosts attribute must be on the same subnet as the other IP addresses for the MultiNICB resource. Type and dimension: string-vector Default: 0.0.0.0 Example: "10.128.8.10, 10.128.8.45"

Table 3-12 Optional attributes

Optional attribute	Description
NoBroadcast	If the value of this attribute is 1, NoBroadcast prevents the agent from sending broadcast pings. ARP requests may still be generated. **Note:** If no NetworkHosts are specified and NoBroadcast is set to 1, the agent cannot function properly. Symantec does not recommend setting NoBroadcast to 1. Type and dimension: integer-scalar Default: 0
OfflineTestRepeatCount	Number of times the test is repeated if the interface status changes from up to down. For every repetition of the test, the next NetworkHosts attribute is selected in round-robin manner. At the end of this process, broadcast is performed if NoBroadcast is set to 0. A greater value prevents spurious changes, but increases the response time. Type and dimension: integer-scalar Default: 3
OnlineTestRepeatCount	The number of times that the test is repeated if the interface changes from down to up. This test helps to prevent oscillations in the status of the interface. Type and dimension: integer-scalar Default: 3
NetworkTimeout	Timeout for ARP and ICMP packets in milliseconds. MultiNICB waits for the response to ICMP and ARP packets only during this time period. Assign the NetworkTimeout a value in the order of tens of milliseconds, given that the ICMP and ARP destinations must be on the local network. Increasing this value increases the time for failover. Type and dimension: integer-scalar Default: 100

Resource type definition

```
type MultiNICB (
    static int OfflineMonitorInterval = 60
    static int MonitorInterval = 10
    static str ArgList[] = { Device, NetworkHosts, Gateway,
    LinkTestRatio, NoBroadcast, NetworkTimeout,
    OnlineTestRepeatCount, OfflineTestRepeatCount }
    static str Operations = None
    str Device{}
    str NetworkHosts[] = { "0.0.0.0" }
    str Gateway
    int LinkTestRatio = 1
    int NoBroadcast
    int NetworkTimeout = 100
    int OnlineTestRepeatCount = 3
    int OfflineTestRepeatCount = 3
)
```

Trigger script

MultiNICB monitor agent function calls a VCS trigger in case of an interface going up or down. The agent passes the following arguments to the script:

- MultiNICB resource name
- The device whose status changed, for example:
  - en0
- The device's previous status (0 for down, 1 for up)
- The device's current status and monitor heartbeat

The agent also sends a notification (which may be received via SNMP or SMTP) to indicate that status of an interface changed. The notification is sent using "health of a cluster resource declined" and "health of a cluster resource improved" traps. These traps are mentioned in the *Veritas Cluster Server User's Guide*. A sample mnicb_postchange trigger is provided with the agent. You can customize this sample script as needed or write one from scratch.

The sample script does the following:

- If interface changes status, it prints a message to the console, for example:

  ```
  MultiNICB agent Res. Name: Device en0 status changed from Down
  to Up.
  ```

Sample configurations

IPMultiNICB and MultiNICB configuration

```
group grp1 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )

    MultiNICB MNICB_grp1 (
        Device@sysa = { en0 = ”10.128.8.42”, en1 = ”10.128.8.43” }
        Device@sysb = { en0 = ”10.128.8.44”, en1 = ”10.128.8.45” }
        NetworkHosts = “10.128.8.10 10.128.8.45”
        LinkTestRatio = 1
        )

    IPMultiNICB ip1 (
        Address = “10.128.10.14”
        Netmask = “255.255.255.0”
        MultiNICBResName = MNICB_grp1
        )
    ip1 requires MNICB_grp1

group grp2 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )

    IPMultiNICB ip2 (
        Address = “10.128.10.15”
        Netmask = “255.255.255.0”
        MultiNICBResName = MNICB_grp1
        )
    Proxy MNICB_proxy (
        TargetResName = MNICB_grp1
        )
    ip2 requires MNICB_proxy
```

DNS agent

The DNS agent updates and monitors the mapping for the following:

- The host name to IP address (A, AAAA, or PTR record)
- The canonical name (CNAME)

The agent performs these tasks for a DNS zone when failing over nodes across subnets (a wide-area failover). Resource records (RR) can include different types: A, AAAA, CNAME, NS (name server), SOA, and PTR records.

Use the DNS agent when the failover source and target nodes are on different subnets. The agent updates the name server and allows clients to connect to the failed over instance of the application service.

For important information about this agent, refer to:

"DNS agent notes" on page 107

Dependencies

No dependencies exist for the DNS resource.

Figure 3-7 Sample service group for a DNS resource

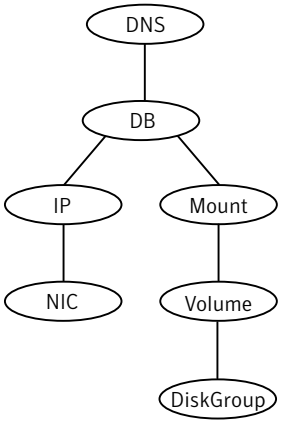

Agent functions

Online	Sends a DNS query to retrieve the Start of Authority (SOA) record of the zone that the Domain agent attribute defines. The master server's name is in the SOA field. Unless you define the StealthMasters attribute, it is the only server for the update. When you define the StealthMasters attribute, only the servers that the attribute defines are updated. The agent creates PTR records for each RR of type A or AAAA if the value of the CreatePTR attribute is true. A prerequisite for this feature is that the same master or stealth servers serve the forward (A or AAAA) and reverse zones.
Offline	Removes the Online lock file. If attribute OffDelRR is true, offline removes all records that the ResRecord keys define.
Monitor	Returns the ONLINE state if at least one name server reports all mappings that ResRecord or Hostname and Alias defines. The name servers are the master or StealthMaster, and all the servers for which an NS record for the zone exists.
Clean	Removes the Online lock file, if it exists.
Open	Removes the Online lock file if the resource is reported online on another node inside the cluster to prevent concurrency violation. If the lock file exists, at least one name server has to report all the RRs that the ResRecord or Hostname and Alias attributes define. If one name server cannot report all the RRs, the agent function removes the Online lock file.
Action	Different action agent functions follow: ■ keyfile.vfd This action entry point checks if the key file as specified in the TSIGKeyFile attribute exists either locally or on shared storage. ■ dig.vfd This action entry point checks if dig and nsupdate binaries exist and are executable. ■ master.vfd This action entry point checks if stealth masters are pingable from the node.

State definitions

ONLINE	Online lock file exists and servers returning all configured resource records.
OFFLINE	Indicates an offline state when either of the following is true: ■ The online lock does not exist. ■ At least one server cannot report all of the RRs' mappings.
UNKNOWN	A problem exists with the configuration. Can indicate that the resource record list contains an invalid value as a part of the record key or a record value of the ResRecord attribute.

Attributes

Table 3-13 Required attributes

Required attribute	Description
Domain	A string representing the DNS zone that the agent administers. The domain name can only contain alphanumeric symbols and the dash. Type and dimension: string-scalar Examples: ■ Forward mapping "demo.symantec.com" ■ IPv4 reverse mapping "2.168.192.in-addr.arpa"
■ Hostname and Alias or ■ ResRecord	You must use either the ResRecord attribute only or the HostName and Alias attributes. Do not use all three attributes together.
Alias	A string representing the alias to the canonical name. Type and dimension: string-scalar Example: "www" Where www is the alias to the canonical name mtv.symantec.com. See “Sample Web server configuration” on page 108.
Hostname	A string that represents the canonical name of a system. Type and dimension: string-scalar Example: "mtv.symantec.com"

Table 3-13 Required attributes

Required attribute	Description
ResRecord	You can use the ResRecord attribute alone, or you can use the Hostname and Alias attributes. ResRecord is an association of DNS resource record values. Each ResRecord attribute consists of two values: *DNS record key* = *DNS record data*. Note that the record key must be a unique value. If the resource record list contains any invalid value as a part of the record key or a record value of the ResRecord attribute, the resource enters an UNKNOWN state. Type and dimension: association-scalar Examples: ■ For forward mapping, where the zone is demo.symantec.com: - sles901 = "192.168.2.191" - ww2 = sles901 - sles9ip6 = "2007::1:2:3:abc" ■ A multi-home DNS record, typically for one host with two network interfaces, different address, but the same DNS name. This results in two-A records, or a single A record with continuation lines. sle902 = "192.168.2.102 10.87.13.22" A multi-home AAAA DNS record can be configured as below: sle902 = "1234::5678 1234::AABB:CCDD" ■ For reverse IPv4 address mapping, where the zone is 2.168.192.in-addr.arpa: 191 = "sles901.demo.symantec.com" ■ For reverse IPv6 address mapping, where the zone is 3.0.0.0.2.0.0.0.1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.7.0.0.2.ip6.arpa: cba = "sles9ip6.demo.symantec.com" Use only partial host names. If you use a fully qualified domain name, append a period "." at the end of the name. For CNAME records, use: ■ ResRecord = { www = mydesktop } or ■ ResRecord = { www = "mydesktop.marketing.db.com." } Where the Domain attribute is "marketing.db.com"

Table 3-14 Required attributes

Required attribute	Description
ResRecord (cont.)	The agent uses case-insensitive pattern matching—and a combination of the Domain and ResRecord attribute values—to determine the resource record type. The RR type is as follows: ■ PTR: if the Domain attribute ends with .arpa ■ A: if the record data field is four sets of numbers, where a space separates each set. The following details the pattern it tries to match: [*1-223*].[*0-255*].[*0-255*].[*0-255*] Hexadecimal is not supported. ■ AAAA: if the record data fields are in multiple sets of hexadecimal format, then this record is an IPv6 associated type AAAA record. ■ CNAME: for any other valid record data. **Note:** If a name in the ResRecord attribute does not comply with RFC 1035, then a warning is issued to the log file. The ResRecord association is not used.

Table 3-15 Optional attributes

Optional attribute	Description
TTL	A non-zero integer represents the “Time To Live” value, in seconds, for the DNS entries in the zone that you want to update. A lower value means more hits on your DNS server, while a higher value means more time for your clients to learn about changes. The time-in-seconds value may take the value 0, which indicates never caching the record, to a maximum of 2,147,483,647, which is over 68 years! The current best practice recommendation (RFC 1912) proposes a value greater than one day, and on RRs that do not change often, consider multi-week values. Type and dimension: integer-scalar Default: 86400 Example: "3600"

Table 3-15 Optional attributes

Optional attribute	Description
StealthMasters	The list of primary master name servers in the domain. This attribute is optional since the first name server is retrieved from the zone's SOA (Start of Authority) record. If the primary master name server is a stealth server, define this attribute. A stealth server is a name server that is authoritative for a zone, but does not appear in that zone's SOA record. It is hidden to prevent direct attacks from the Internet. Type and dimension: string-keylist Example: { "10.190.112.23" }
TSIGKeyFile	Required when you configure DNS for secure updates. Specifies the absolute path to the file containing the private TSIG (Transaction Signature) key. Type and dimension: string-scalar Example: /var/tsig/example.com.+157+00000.private
CreatePTR	Use the CreatePTR attribute to direct the online agent function to create PTR records for each RR of type A or AAAA. You must set the value of this attribute to true (1) to create the records. Before you can use this attribute, the same master or stealth servers must serve the forward (A or AAAA) and reverse zones. Type and dimension: boolean-scalar Default: 0 Example: 1
OffDelRR	Use the OffDelRR attribute to direct the offline agent function to remove all records that the ResRecord key defines. You must set the value of this attribute to true (1) to have the agent remove all the records. The online agent function always adds records if they do not exist. Type and dimension: boolean-scalar Default: 0 Example: 1

Resource type definition

```
type DNS (
    static keylist SupportedActions = { "dig.vfd", "keyfile.vfd",
    "master.vfd" }
    static str ArgList[] = { Domain, Alias, Hostname, TTL,
    TSIGKeyFile, StealthMasters, ResRecord, CreatePTR, OffDelRR }
    str Domain
    str Alias
    str Hostname
    int TTL = 86400
    str StealthMasters[]
    str TSIGKeyFile
    str ResRecord{}
    boolean CreatePTR = 0
    boolean OffDelRR = 0
)
```

DNS agent notes

The DNS agent has the following notes:

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node; discrepancies that might prevent a service group from going online on a specific node.

For DNS resources, the high availability drill tests the following conditions:

- Checks if the key file as specified by the TSIGKeyFile attribute is available either locally or on shared storage.
- Checks if the dig and nsupdate binaries are available on the cluster node and are executable on that node.
- Checks if the stealth masters are pingable from the cluster node so as to ensure that there is no network issue that would prohibit the DNS update and query requests from reaching the stealth master server.

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Monitor scenarios

Depending on the existence of the Online lock file and the defined Resource Records (RR), you get different status messages from the Monitor function.

Table 3-16 Monitor scenarios for the Online lock file

Online lock file exists	Expected RR mapping	Monitor returns
NO	N/A	OFFLINE
YES	NO	OFFLINE
YES	YES	ONLINE

Sample Web server configuration

Take the former Veritas corporate web server as an example. A browser requests the URL http://www.veritas.com that maps to the canonical name mtv.veritas.com. The browser retrieves the IP address for the web server by querying a domain name server. If the web server fails over from Mountain View to Heathrow (hro.veritas.com), the domain name servers need a new canonical name mapping for www.veritas.com. The www.veritas.com alias is now updated to point to the canonical name of the standby system in Heathrow.

Secure DNS update for BIND 9

The DNS agent expects that the zone's allow-update field contains the IP address for the hosts that can dynamically update the DNS records. This functionality is default for the DNS agent. Since a competent black hat can, however, spoof IP addresses, consider TSIG as an alternative.

TSIG (Transaction Signature) as specified in RFC 2845 is a shared key message authentication mechanism that is available in DNS. A TSIG key provides the means to authenticate and verify the validity of exchanged DNS data. It uses a shared secret key between a resolver and either one or two servers to provide security.

Setting up secure updates using TSIG keys for BIND 9

In the following example, the domain is example.com.

To use secure updates using TSIG keys

1 Run the `dnssec-keygen` command with the HMAC-MD5 option to generate a pair of files that contain the TSIG key:

```
# dnssec-keygen -a HMAC-MD5 -n HOST example.com.
    example.com.+157+00000
```

2 Open the example.com.+157+00000.key file. After you run the `cat` command, the contents of the file resembles:

```
# cat example.com.+157+00000.key
    example.com. IN KEY 512 3 157 +Cdjlkef9ZTSeixERZ433Q==
```

3 Copy the shared secret (the TSIG key), which looks like:

```
+Cdjlkef9ZTSeixERZ433Q==
```

4 Configure the DNS server to only allow TSIG updates using the generated key. Open the named.conf file and add these lines.

```
key example.com. {
    algorithm hmac-md5;
    secret "+Cdjlkef9ZTSeixERZ433Q==";
};
```

Where **+Cdjlkef9ZTSeixERZ433Q==** is the key.

5 In the named.conf file, edit the appropriate zone section and add the allow-updates sub-statement to reference the key:

```
allow-update { key example.com. ; } ;
```

6 Save and restart the named process.

7 Place the files containing the keys on each of the nodes that is listed in your group's SystemList. The DNS agent uses this key to update the name server. Copy both the private and public key files on to the node. A good location is in the /var/tsig/ directory.

8 Set the TSIGKeyFile attribute for the DNS resource to specify the file containing the private key.

```
DNS www (
Domain = "example.com"
ResRecord = {www = north}
TSIGKeyFile = "/var/tsig/example.com.+157+00000.private"
)
```

Chapter 4

File share agents

This chapter contains the following:

About the file service agents

Use the file service agents to provide high availability for file share resources.

NFS agent

Starts and monitors the nfsd and mountd subsystem processes required by all exported NFS file systems.

The srcmstr daemon is the System Resource Controller (SRC). This agent sends requests to the SRC to start and monitor these daemons. You need to start the srcmstr daemon before using this agent.

Note: NFSv4root and NFSSecurity require AIX 5.3 ML03.

Symantec recommends that you configure only one NFS resource in a service group on a node. If you have more than one service group that uses the NFS resource, the other service groups can use a Proxy resource. The Proxy resource can point to the NFS resource in the first group. This use of the Proxy resource prevents redundant monitoring of the NFS daemons on the same system.

For important information about this agent, refer to:

"NFS agent notes" on page 115

Dependencies

The NFS resource does not depend on other resources.

Figure 4-1 Sample service group for an NFS resource

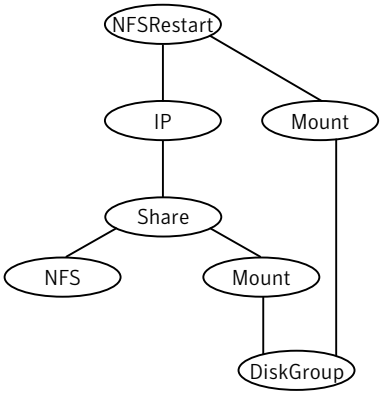

Agent functions

Online	■ Checks if nfsd and mountd are running. If they are not running, the agent starts the daemons and exits. ■ The nfsrgyd daemon is started if NFSv4Root is specified. ■ The gssd daemon is started if NFSSecurity is set to 1.
Offline	Not applicable.
Monitor	■ Monitors nfsd and mountd by checking whether or not the daemons are active. ■ The nfsrgyd daemon is monitored if NFSv4Root is specified. ■ The gssd daemon monitored if NFSSecurity is set to 1.
Clean	Terminates the resource and takes it offline—forcibly if necessary.

State definitions

ONLINE	Indicates that the NFS daemons are running in accordance with the supported protocols and versions.
OFFLINE	Indicates that the NFS daemons are not running in accordance with the supported protocols and versions.
FAULTED	Indicates that the NFS daemons are not running in accordance with the supported protocols and versions.
UNKNOWN	Unable to determine the status of the NFS daemons.

Attributes

Table 4-1 Optional attributes

Optional attribute	Description
Nservers	Specifies the number of concurrent NFS requests the server can handle. Type and dimension: integer-scalar Default: 10

Table 4-1 Optional attributes

Optional attribute	Description
NFSv4Root	Root directory of the NFSv4 pseudo file system to be exported. All exports should have a path relative to the path specified by this attribute. You can explicitly create the NFSv4 pseudo file system by specifying the exname option of the `exportfs` command in the Options attribute of the Share resource. If you want to export file systems with NFSv4 protocols and do not want to explicitly create NFSv4 pseudo file system by using the exname option, then set NFSv4Root to "/". Required for filesystems to be exported with v4 protocol. Type and dimension: string-scalar
NFSSecurity	If the value of this attribute is 1,the gssd daemon starts. You must configure the type of security that NFS supports, for example: Kerberos. Type and dimension: boolean-scalar Default: 0
GracePeriod	Specifies the grace period for which the server allows lock recovery. Required for NFS lock recovery. Type and dimension: integer-scalar
LockFileTimeout	Specify the amount of time required, in seconds, for the service group to go online. The agent uses this attribute to synchronize the starting and stopping of daemons between multiple service groups. Type and dimension: integer-scalar Default: 180 Example: "240""

Resource type definition

```
"type NFS (
    static int RestartLimit = 1
    static str ArgList[] = { Nservers, GracePeriod, NFSv4Root,
    NFSSecurity, LockFileTimeout }
    static str Operations = OnOnly
    int Nservers = 10
    int GracePeriod = 90
    str NFSv4Root
    boolean NFSSecurity = 0
    int LockFileTimeout = 180
)
```

NFS agent notes

The NFS agent has the following notes:

- "Caveat when nodes in the cluster are a mix of AIX operating system versions 5.x and 6.1" on page 115
- "Using NFSv4" on page 115

Caveat when nodes in the cluster are a mix of AIX operating system versions 5.x and 6.1

Failing over an exported file system between the NFS server nodes with different AIX operating system versions can result in a Stale file handle error at the NFS client. This issue is independent of VCS.

Using NFSv4

For NFS v4 support, you must specify the NFSv4Root attribute. You must include `vers=4` in the Option attribute of the Share resource.

Set up Enterprise Identity Mapping (EIM) in the NFS environment, if:

- Mapping of userids and username is not same on both client and server
- Client and server belong to different domains

If either of the above points are true, and EIM is not set up, the client has minimal rights (user=nobody, group=nobody).

If you want to use the NFSv4 security feature, set the NFSSecurity attribute of the NFS resource to `1`. Manually configure Kerberos or any other security environment that is supported by NFSv4.

Caveats

You export filesystems with NFSv4Root="/exp/exports1", and you forcefully stop the engine so that exports are still valid and existing. If you change configurations on NFS to set NFSv4Root="/newexport", the NFS Agent is not able to come online with this new root, because the already exported filesystem is using an older NFS pseudo file system root. To avoid this problem bring all Share resources down properly before changing NFSv4Root.

If you create a pseudo file system, a client can access the filesystem. After the NFS server fails over to the other system in the cluster, the client can not see the filesystem. The client needs to remount it.

Sample configurations

Configuration 1

```
include "types.cf"

cluster vcs_cluster (
        CounterInterval = 5
        )

system sysa (
        )

system sysb (
        )

group test_grp (
        SystemList = { sysa = 0, sysb = 1 }
        )

        DiskGroup test_dg (
                DiskGroup = test_dg
                )

        IP test_ip (
                Device = en0
                Address = "10.182.13.28"
                NetMask = "255.255.240.0"
                )

        Mount test_mnt (
                MountPoint = "/test_mnt"
                BlockDevice = "/dev/vx/dsk/test_dg/test_vol"
                FSType = vxfs
                MountOpt = rw
                FsckOpt = "-y"
                )
```

```
NFS test_nfs (
        Nservers = 20
)

NFSRestart test_nfsrestart (
)

Share test_share (
        PathName = "/test_mnt"
        Options = "-o rw,root=vcsaix3"
        )

Volume test_vol (
        Volume = test_vol
        DiskGroup = test_dg
        )


test_nfsrestart requires test_ip
test_ip requires test_share
test_share requires test_nfs
test_share requires test_mnt
test_mnt requires test_vol
test_vol requires test_dg
```

Configuration 2

This is a sample VCS configuration for using the NFSv4 feature by explicitly creating the NFSv4 pseudo file system. Here, NFSv4Root is set to "/export", which is different than the root of the local file system. If you want to use NFSv4 features and the NFSv4Root attribute is not equal to "/", you must configure the Share resources with exname in the Options attribute relative to the NFSv4Root. In this example it is: exname="/export/export1".

```
include "types.cf"

cluster vcs_cluster (
        CounterInterval = 5
        )

system sysa (
        )

system sysb (
        )

group test_grp (
        SystemList = { sysa = 0, sysb = 1 }
        )
```

```
DiskGroup test_dg (
        DiskGroup = test_dg
        )

IP test_ip (
        Device = en0
        Address = "10.182.13.28"
        NetMask = "255.255.240.0"
        )

Mount test_mnt (
        MountPoint = "/test_mnt"
        BlockDevice = "/dev/vx/dsk/test_dg/test_vol"
        FSType = vxfs
        MountOpt = rw
        FsckOpt = "-y"
        )

NFS test_nfs (
        Nservers = 20
        NFSv4Root = "/export"
)

NFSRestart test_nfsrestart (
)

Share test_share (
        PathName = "/test_mnt"
        Options = "-o rw,vers=4,root=vcsaix3,exname=/ \
        export/export1"
        )

Volume test_vol (
        Volume = test_vol
        DiskGroup = test_dg
        )


test_nfsrestart requires test_ip
test_ip requires test_share
test_share requires test_nfs
test_share requires test_mnt
test_mnt requires test_vol
test_vol requires test_dg
```

NFSRestart agent

The NFSRestart agent recovers NFS record locks after sudden reboots or crashes on clients and servers. This avoids file corruption and provides high availability for NFS record locks.

The NFSRestart agent brings online, takes offline, and monitors the three daemons: smsyncd, statd, and lockd. If you have configured the NFSRestart agent for lock recovery, the NFSRestart agent starts the smsyncd daemon. The daemon copies the NFS information on the location of connections from the shared-storage to the local directory (/var/statmon/sm) and vice-versa.

Note: NFSv4root and NFSSecurity require AIX 5.3 ML03.

For important information about this agent, refer to:

"NFSRestart agent notes" on page 122

Dependencies

This resource must be at the top of the resource dependency tree of a service group. Only one NFSRestart resource should be configured in a service group. The NFSRestart, NFS, and Share agents must be in same service group.

Figure 4-2 Sample service group for an NFSRestart resource

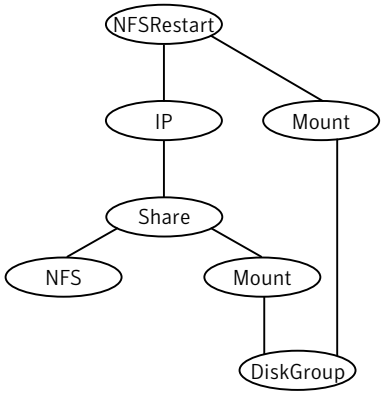

Agent functions

Online	■ Terminates statd and lockd. ■ If the value of the NFSLockFailover attribute is 1, it copies the locks from the shared storage to the /var/statmon/sm directory. ■ Copies the locks from the shared storage to the /var/statmon/sm directory if NFSLockFailover is set to 1. ■ Starts the statd and lockd daemons. ■ Starts the smsyncd daemon to copy the contents of the /var/statmon/sm directory to the shared storage (LocksPathName) at regular, two-second intervals if the value of the NFSLockFailover attribute is 1.
Monitor	Monitors the statd and lockd daemons and restarts them if they are not running. It also monitors the smsyncd daemon if the value of the NFSLockFailover attribute is 1.
Offline	■ Terminates the statd and lockd daemons to clear the lock state. ■ Terminates the nfsd and mountd daemons to close the TCP/IP connections. ■ Terminates the smsyncd daemon if the daemon is running.
Clean	■ Terminates the statd and lockd daemons to clear the lock state. ■ Terminates the nfsd and mountd daemons to close TCP/IP connections. ■ Terminates the smsyncd daemon if the daemon is running.
nfs_postoffline	■ Restarts nfsd, mountd, lockd and statd after the group goes offline. ■ The nfsrgyd daemon is restarted if NFSv4Root is set. ■ The gssd daemon is restarted if NFSSecurity is set to 1.
Action	■ nfsconf.vfd Checks the NFS configuration file to confirm that the NFS server does not come online automatically after reboot. ■ lockdir.vfd Verifies that the NFS lock directory (which is specified by the LocksPathName attribute of NFSRestart) is on shared storage.

State definitions

ONLINE	Indicates that the daemons are running properly.
OFFLINE	Indicates that one or more daemons are not running.
UNKNOWN	Indicates the inability to determine the agent's status.

Attributes

Table 4-2 Optional attributes

Optional attribute	Description
LocksPathName	The path name of the directory to store the NFS lock information. This attribute is required when the value of the NFSLockFailover attribute is 1. The path that you specify for the LocksPathName attribute should be on shared storage. This is to ensure that it is accessible to all the systems where the NFSRestart resource fails over. Type and dimension: string-scalar
NFSLockFailover	NFS Lock recovery is done for all the Share resources that are configured in the group of this resource. Type and dimension: boolean-scalar Default: 0
NFSRes	Name of the NFS resource on the system. This attribute is required if the value of the NFSLockFailover attribute is 1. Type and dimension: string-scalar

Resource type definition

```
type NFSRestart (
    static keylist SupportedActions = { "lockdir.vfd",
    "nfsconf.vfd" }
    static str ArgList[] = { NFSLockFailover, LocksPathName,
    "NFSRes:GracePeriod", "NFSRes:LockFileTimeout"  }
    str LocksPathName
    str NFSRes
    boolean NFSLockFailover = 0
)
```

NFSRestart agent notes

The NFSRestart agent has the following notes:

- "High availability fire drill" on page 122
- "Providing a fully qualified host name" on page 122
- "Providing a fully qualified host name" on page 123

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node; discrepancies that might prevent a service group from going online on a specific node. For NFSRestart resources, the high availability drill performs the following, it:

- Checks the NFS configuration file to confirm that the NFS server does not come online automatically after reboot.
- Verifies that the NFS lock directory (which is specified by the LocksPathName attribute of NFSRestart) is on shared storage.

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Providing a fully qualified host name

You must provide a fully qualified host name (nfsserver.princeton.edu) for the NFS server while mounting the file system on the NFS client. If you do not use a fully qualified host name, or if you use a virtual IP address (10.122.12.25) or partial host name (nfsserver), NFS lock recovery fails.

If you want to use the virtual IP address or a partial host name, make the following changes to the service database (hosts) and the netsvc.conf files:

```
/etc/hosts
```

To use the virtual IP address and partial host name for the NFS server, you need to add an entry to the /etc/hosts file. The virtual IP address and the partial host name should resolve to the fully qualified host name.

```
/etc/netsvc.conf
```

You should also modify the hosts entry in this file so that upon resolving a name locally, the host does not first contact NIS/DNS, but instead immediately returns a successful status. Changing the netsvc.conf file might affect other services running on the system.

For example:

```
hosts = local,bind,nis
```

You have to make sure that the NFS client stores the same information for the NFS server as the client uses while mounting the file system. For example, if the NFS client mounts the file system using fully qualified domain names for the NFS server, then the NFS client directory: /var/statmon/sm directory should also have a fully qualified domain name after the acquisition of locks. Otherwise, you need to start and stop the NFS client twice using the /etc/init.d/nfs.client script to clear the lock cache of the NFS client.

A time period exists where the virtual IP address is online but locking services are not registered on the server. Any NFS client trying to acquire a lock in this interval would fail and get ENOLCK error.

Every two seconds, the smsyncd daemon copies the list of clients that hold the locks on the shared filesystem in the service group. If the service group fails before smsyncd has a chance to copy the client list, the clients may not get a notification once the service group is brought up. This causes NFS lock recovery failure.

Providing a fully qualified host name

You must provide a fully qualified host name (nfsserver.princeton.edu) for the NFS server while mounting the file system on the NFS client. If you do not use a fully qualified host name, or if you use a virtual IP address (10.122.12.25) or partial host name (nfsserver), NFS lock recovery fails.

If you want to use the virtual IP address or a partial host name, make the following changes to the service database (hosts) and the nsswitch.conf files:

```
/etc/hosts
```

To use the virtual IP address and partial host name for the NFS server, you need to add an entry to the /etc/hosts file. The virtual IP address and the partial host name should resolve to the fully qualified host name.

```
/etc/nsswitch.conf
```

You should also modify the hosts entry in this file so that upon resolving a name locally, the host does not first contact NIS/DNS, but instead immediately

returns a successful status. Changing the nsswitch.conf file might affect other services running on the system.

For example:

```
hosts:  files [SUCCESS=return] dns nis
```

You have to make sure that the NFS client stores the same information for the NFS server as the client uses while mounting the file system. For example, if the NFS client mounts the file system using fully qualified domain names for the NFS server, then the NFS client directory: /var/statmon/sm directory should also have a fully qualified domain name after the acquisition of locks. Otherwise, you need to start and stop the NFS client twice using the /etc/init.d/nfs.client script to clear the lock cache of the NFS client.

A time period exists where the virtual IP address is online but locking services are not registered on the server. Any NFS client trying to acquire a lock in this interval would fail and get ENOLCK error.

Every two seconds, the smsyncd daemon copies the list of clients that hold the locks on the shared filesystem in the service group. If the service group fails before smsyncd has a chance to copy the client list, the clients may not get a notification once the service group is brought up. This causes NFS lock recovery failure.

Sample configurations

```
include "types.cf"

cluster vcs_cluster (
    CounterInterval = 5
    )

system sysa (
    )

system sysb (
    )

group test_grp (
    SystemList = { sysa = 0, sysb = 1 }
    )

    DiskGroup test_dg (
        DiskGroup = test_dg
        )

    IP test_ip (
        Device = en0
        Address = "10.182.13.28"
        NetMask = "255.255.240.0"
        )
```

```
Mount test_mnt (
    MountPoint = "/test_mnt"
    BlockDevice = "/dev/vx/dsk/test_dg/test_vol"
    FSType = vxfs
    MountOpt = rw
    FsckOpt = "-y %-o full"
    )

Mount test_lockinfo_mnt (
    MountPoint = "/lockinfo"
    BlockDevice = "/dev/vx/dsk/test_dg/test_lockinfo_vol"
    FSType = vxfs
    MountOpt = rw
    FsckOpt = "-y"
    )

NFS test_nfs (
    Nservers = 20
)

NFSRestart test_nfsrestart (
    NFSLockFailover  = 1
    LocksPathName = "/test_mnt"
    NFSRes = test_nfs
    )

Share test_share (
    PathName = "/test_mnt"
    Options = "-o rw"
    )

Volume test_lockinfo_vol (
    Volume = test_lockinfo_vol
    DiskGroup = test_dg
    )

Volume test_vol (
    Volume = test_vol
    DiskGroup = test_dg
    )

test_nfsrestart requires test_ip
test_nfsrestart requires test_lockinfo_mnt
test_lockinfo_mnt requires test_lockinfo_vol
test_lockinfo_vol requires test_dg
test_ip requires test_share
test_share requires test_nfs
test_share requires test_mnt
test_mnt requires test_vol
test_vol requires test_dg
```

Share agent

Shares, unshares, and monitors a single local resource for exporting an NFS file system to be mounted by remote systems.

Before you use this agent, verify that the files and directories to be shared are on shared disks.

For important information on this agent, refer to:

"Share agent notes" on page 128

Dependencies

Share resources depend on NFS. In NFS service group, IP, IPMultiNIC, and IPMultiNICB resources depend on Share resources.

Figure 4-3 Sample service group for a Share resource

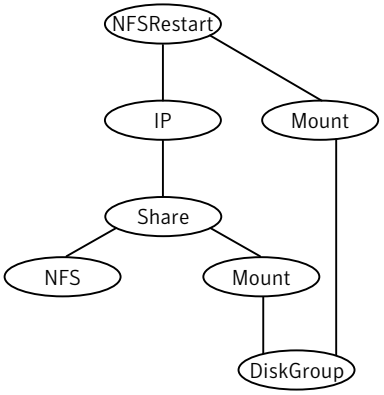

Agent functions

Online	Exports (shares) a directory to the specified client.
Offline	Unshares the exported directory from the client.
Monitor	Verifies that the shared directory is exported to the client.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

Action	direxists.vfd Checks if the path specified by the PathName attribute exists on the cluster node. If the path name is not specified, it checks if a corresponding mount point is available to ensure that the path is on shared storage.

State definitions

ONLINE	Indicates that specified directory is exported to the client.
OFFLINE	Indicates that the specified directory is not exported to the client.
UNKNOWN	Indicates that the agent could not determine the state of the resource or that the resource attributes are invalid.
FAULTED	Indicates that the share has unexported outside of VCS control.

Attributes

Table 4-3 Required attributes

Required attribute	Description
PathName	Pathname of the file system to be shared. Type and dimension: string-scalar Example: "/share1x"

Table 4-4 Optional attributes

Optional attribute	Description
Options	Options to the `exportfs` command. When specifying multiple options, separate them with commas, for example: "rw,vers=4" For more information about the `exportfs` command and its options, refer to the exportfs manpage. Type and dimension: string-vector

Resource type definition

```
type Share (
    static keylist SupportedActions = { "direxists.vfd" }
    static str ArgList[] = { PathName, Options }
    str PathName
    str Options
)
```

Share agent notes

The following section contains notes on the Share agent.

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node. These discrepancies might prevent a service group from going online on a specific node. For Share resources, the high availability fire drill checks if the path exists.

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Sample configurations

Configuration

The options for the Share agent must be specific. In this configuration, root access is only given to the system sysA. By default, systems do not have root access to the exported directory /mnt1. Refer to the exportfs manual page for more information.

```
Share share1 (
    PathName = "/mnt1"
    Options = "rw=sysA,access=sysA,root=sysA"
    )
```

About the Samba agents

Samba is a suite of programs that allows a system running a UNIX or UNIX-like operating system to provide services using the Microsoft network protocol. Samba supports the following services:

- Filespace
- Printer
- WINS
- Domain Master

Configure these services in the Samba configuration file (smb.conf). Samba uses two processes: smbd and nmbd to provide these services.

VCS provides Samba failover using three agents: SambaServer, NetBios, and SambaShare.

The Samba agents

- The NetBios agent
- The SambaServer agent
- The SambaShare agent

Before using the Samba agents

- Verify that smbd and nmbd always run as daemons. Verify that they cannot be started using the meta-daemon inetd.
- Verify that the smbd and nmbd daemons are in the path environment variable.
- If they are not, verify that they run from the default directory /usr/bin.
  - The path of smbd and nmbd is /usr/local/samba/sbin.
- Verify that Samba is configured properly and that the Samba configuration file is identical on all cluster systems. The user can replicate the file or store it on a shared disk accessible from all cluster systems.
- If configuring Samba as a WINS server or Domain Master, verify that the Samba lock directory is on the shared disk. This ensures that the WINS server database and Domain Master are created on the shared disk.

Supported versions

Table 4-5 provides the support matrix for the Samba agents.

Table 4-5 Supported platforms, architectures, and Samba versions

Platforms	Operating systems/ Architecture	Supported Samba versions
AIX	5.3/6.1	Version 3.0.24 or later

Configuring the Samba agents

If Samba is configured properly, and the configuration file is identical on all cluster systems, configure resources of type SambaServer and NetBios only. This ensures that all shares in the Samba configuration file are failed over when the SambaServer resource fails over. Note that the Samba shares are not monitored. To monitor the Samba shares, configure the agents with the following dependencies:

```
SambaShare requires NetBios
SambaShare requires SambaServer
NetBios requies IP
```

For example, use the following configuration to monitor Samba shares SambaShare1 and SambaShare2. Use multiple resources of type SambaShare (if necessary), but only one resource each of type NetBios and SambaServer.

```
SambaShare1 requires NetBios1
SambaShare1 requires SambaServer1
SambaShare2 requires NetBios1
SambaShare2 requires SambaServer1
NetBios1 requies IP_1
```

SambaServer agent

The SambServer agent starts, stops, and monitors the smbd process as a daemon. Only one resource of this type is permitted. You can use the agent to make a smbd daemon highly available or to monitor it.

The smbd daemon provides Samba share services. The agent makes a copy of smbd for each client and verifies that Samba is running by reading the pid of this daemon. The agent can perform in-depth monitoring by establishing a socket connection to Samba at ports where the daemon is listening and sending it a NetBIOS session request.

Dependencies

No dependencies exist for the SambaServer resource.

Figure 4-4 Sample service group for a SambaServer resource

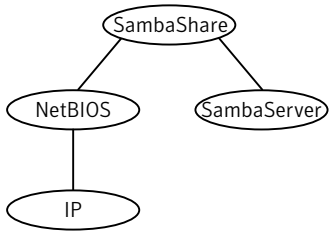

Agent functions

Online	Starts the smbd daemon at specified ports.
Offline	Stops the smbd daemon.
Monitor	Verifies that the smbd daemon is running by reading its pid file. Does in-depth monitoring periodically, if configured, by establishing a socket connection to Samba and sending it a NetBIOS session request.
Clean	Stops the smbd daemon.

State definitions

ONLINE	Indicates that the smbd daemon is running. If in-depth monitoring is configured, it indicates that a positive session response packet was received through a socket connection to the Samba server.
OFFLINE	Indicates that smbd is not running. If in-depth monitoring is enabled, it indicates that the agent could not establish a socket connection with the server, or that it received an incorrect response packet header, or the session response packet connection timed out.
UNKNOWN	Indicates that the agent could not determine the state of the resource.

Attributes

Table 4-6 Required attributes

Required attribute	Description
ConfFile	Complete path of the configuration file that Samba uses. Type and dimension: string-scalar Example: "/etc/sfw/smb.conf"
LockDir	Lock directory of Samba. Samba stores the files smbd.pid, nmbd.pid, wins.dat (WINS database), and browse.dat (master browser database) in this directory. Type and dimension: string-scalar Example: “ /usr/local/samba/var/locks”
SambaTopDir	Parent path of Samba daemon and binaries. Example: “/usr/local/samba”

Resource type definitions

```
type SambaServer (
    static int RestartLimit = 5
    static str ArgList [] = {ConfF, LockDir,
    IndepthMonitorCyclePeriod}
    str ConfFile = “etc/smb.conf”
    str LockDir = "/usr/local/samba/var/locks"
    str SambaTopDir = "/usr/local/samba"
    str LockDir = “/var/lock/samba”
    int IndepthMonitorCyclePeriod = 5
    int ResponseTimeout = 10
)
```

Sample configurations

```
SambaServer samba_server (
    ConfFile = "/etc/smb.conf"
    LockDir = "/usr/local/samba/var/locks"
    SambaTopDir = "/usr/local/samba"
    IndepthMonitorCyclePeriod = 3
    ResponseTimeout = 15
)
```

SambaShare agent

The SambaShare agent adds, removes, and monitors a share by modifying the specified Samba configuration file. You can use the agent to make a Samba Share highly available or to monitor it.

Each filespace or printer service provided by Samba is a shared resource and is defined as a section in the Samba configuration file. The section name is the name of the shared resource and the section parameters define the share attributes.

Dependencies

SambaShare resources depend on SambaServer, NetBios and Mount resources.

Figure 4-5 Sample service group for a SambaShare resource

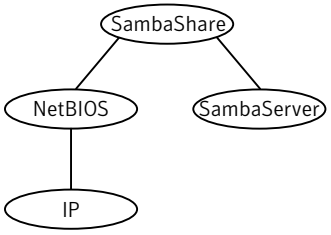

Agent functions

Online	Edits the samba configuration file and adds the shares.
Offline	Removes the shares from the configuration file.
Monitor	Issues the command `smbclient` to check if the specified shares exist.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	Indicates that the share is available and that the share path exists.
OFFLINE	Indicates that the share is not available, or that the share has a non-existent path.
UNKNOWN	Indicates that the agent could not determine the state of the resource.

Attributes

Table 4-7 Required attributes

Required attribute	Description
SambaServerRes	Name of the SambaServer resource. Type and dimension: string-scalar Example: "SG.smb_res1" Where SG is the service group to which the resource smb_res1 belongs.
"SambaServerRes:ConfFile"	Complete path of the configuration file that is specified in the SambaServer resource.
"SambaServerRes:LockDir"	Complete path of the lock directory that is specified in the SambaServer resource.
"SambaServerRes:SambaTopDir"	Parent path of Samba daemon and binaries installed.
ShareName	Name of the share resource. Type and dimension: string-scalar Example: "share1"
ShareOptions	List of parameters for the share attributes. These parameters are specified as name=value pairs, with each pair separated by a semicolon (;). Type and dimension: string-scalar Example: "path=/shared; public=yes; writable=yes"

Resource type definition

```
type SambaShare (
    static str ArgList[] = { "SambaServerRes:ConfFile",
    "SambaServerRes:SambaTopDir", "SambaServerRes:LockDir",
    ShareName, ShareOptions, "SambaServerRes:Ports" }
    str SambaServerRes
    str ShareName
    str ShareOptions
)
```

Sample configuration

```
SambaShare Samba_SambaShare3 (
        SambaServerRes = Samba_SambaServer
        ShareName = smbshare3
        ShareOptions = "path=/smbshare3; public=yes; writable=yes"
        )
```

NetBIOS agent

The NetBIOS agent starts, stops, and monitors the nmbd daemon. Only one resource of this type is permitted. You can use the agent to make the nmbd daemon highly available or to monitor it.

The agent sets, monitors, and resets the names and network interfaces by which the Samba server is known. The agent also sets, monitors and resets Samba to act as a WINS server or domain master or both.

Note that nmbd broadcasts the NetBIOS name, or the name by which the Samba server is known in the network.

Before using this agent:

- Set the NetBIOS name.
- Set the NetBIOS interface.

Dependencies

The NetBios resource depends on the IP, the IPMultiNIC or the IPMultiNICB resource.

Note: You can configure only one NetBios resource on a system.

Figure 4-6 Sample service group for a NetBIOS resource

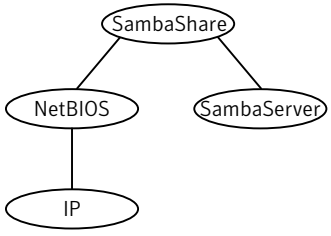

Agent functions

Online	Updates the Samba configuration with the NetBIOS name, all NetBIOS aliases and network interfaces, WINS support, and domain master options specified in the NetBIOS resource. Starts the nmbd daemon.
Offline	Removes the NetBIOS name, all NetBIOS aliases and network interfaces, WINS support, and domain master options specified in the NetBIOS resource from the Samba configuration file. Stops the nmbd daemon.
Monitor	Verifies that the Samba configuration contains the NetBIOS name, all NetBIOS aliases and network interfaces, WINS support, and domain master options specified in the NetBIOS resource.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	Indicates that the specified NetBIOS aliases are advertised and that Samba is handling requests for all specified network interfaces. Indicates that WINS and Domain support services are running, if configured.
OFFLINE	Indicates one or more of the following: ■ NetBIOS name is not advertised. ■ A NetBIOS alias is not advertised. ■ Samba is not handling requests on one of the specified interfaces. ■ If WINS support is configured, Samba is not providing WINS service. ■ If domain support is set, Samba is not providing Domain Master service.
UNKNOWN	Indicates that the agent could not determine the state of the resource.

Attributes

Table 4-8 Required attributes

Required attribute	Description
NetBiosName	Name by which the Samba server is known in the network. Type and dimension: string-scalar
"SambaServerRes:ConfFile"	Complete path of the configuration file that is specified in the SambaServer resource. Type and dimension: string-scalar
"SambaServerRes:LockDir"	Complete path of the lock directory that is specified in the SambaServer resource.
SambaServerRes:SambaTopDir	Parent path of Samba daemon and binaries installed.

Table 4-9 Optional attributes

Optional attribute	Description
Interfaces	List of network interfaces on which Samba handles browsing. Type and dimension: string-vector Example: "172.29.9.24/16"
NetBiosAliases	List of additional names by which the Samba server is known in the network. Type and dimension: string-vector Example: "host1_samba, myname"

Table 4-9 Optional attributes

Optional attribute	Description
WinsSupport	If set to 1, this flag causes the agent to configure Samba as a WINS server. Type and dimension: integer-scalar Default: 0

Resource type definition

```
type NetBios (
    static str ArgList[] = { "SambaServerRes:ConfFile",
    "SambaServerRes:SambaTopDir", "SambaServerRes:LockDir",
    NetBiosName, NetBiosAliases, Interfaces, WinsSupport,
    DomainMaster }
    str SambaServerRes
    str NetBiosName
    str NetBiosAliases[]
    str Interfaces[]
    int WinsSupport
    int DomainMaster
)
```

Sample configuration

```
NetBios Samba_NetBios (
    SambaServerRes = Samba_SambaServer
    NetBiosName = samba_demon
    NetBiosAliases = { asamba_demon, samba127 }
    WinsSupport = 1
    DomainMaster = 1
    )
```

Chapter 5

Service and application agents

This chapter contains the following agents:

About the service and application agents

Use service and application agents to provide high availability for application and process-related resources.

Apache Web server agent

The Apache Web server agent brings an Apache Server online, takes it offline, and monitors its processes. The Apache Web server agent consists of resource type declarations and agent scripts. You use the Apache Web server agent, in conjunction with other agents, to make an Apache Web server highly available.

This agent supports the Apache HTTP server 1.3, 2.0, and 2.2. It also supports the IBM HTTP Server 1.3 and 2.0.

This agent can detect when an Apache Web server is brought down gracefully by an administrator. When Apache is brought down gracefully, the agent does not trigger a resource fault even though Apache is down.

Note: The Apache agent requires an IP resource for operation.

For more information regarding this agent:

See "Apache Web server notes" on page 149.

Dependencies

This type of resource depends on IP and Mount resources.

Figure 5-1 Sample service group for the Apache Web server agent

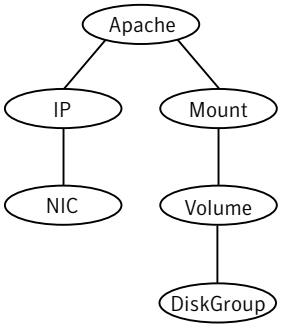

Agent functions

Online	Starts an Apache server by executing the httpdDir/httpd program with the appropriate arguments. When you specify a file with the EnvFile attribute, the file is sourced before the agent executes the httpd command.
Offline	To stop the Apache HTTP server, the agent: ■ Executes the httpdDir/httpd program with the appropriate arguments (Apache v2.0), or ■ Sends a TERM signal to the HTTP Server parent process (Apache v1.3). When you specify a file with the EnvFile attribute, the file is sourced before the agent executes the httpd command.
Monitor	Monitors the state of the Apache server. First it checks for the processes, next it can perform an optional state check.
Clean	Removes the Apache HTTP server system resources that might remain after a server fault or after an unsuccessful attempt to online or offline. These resources include the parent httpd daemon and its child daemons.
Action	checkconffile.vfd Checks for the existence of the Apache configuration file and the existence of the directory that contains the httpd binary that is used during start up. For a local installation, if the config file or HttpdDir is not found, make sure that it exists on the failover node.

State definitions

ONLINE	Indicates that the Apache server is running.
OFFLINE	Indicates that the Apache server is not running. Can also indicate that the administrator has stopped the Web server gracefully. Note that the agent uses the PidFile attribute for intentional offline detection.
UNKNOWN	Indicates that a problem exists with the configuration.

Attributes

Table 5-1 Required attributes

Required attribute	Description
ConfigFile	Full path and file name of the main configuration file for the Apache server. Type and dimension: string-scalar Example: "/apache/server1/conf/httpd.conf"
httpdDir	Full path of the directory to the httpd binary file Type and dimension: string-scalar Example: "/apache/server1/bin"
SecondLevelMonitor	Enables second-level monitoring for the resource. Second-level monitoring is a deeper, more thorough state check of the Apache HTTP server. An HTTP GET request on the Web server's root directory performs the monitoring. Valid attribute values are 1 (true) and 0 (false). Specifying this attribute is required. Type and dimension: boolean-scalar Default: 0 Example: "1"
ResLogLevel	Controls the agent's logging detail for a specific instance of a resource. Values are: ■ ERROR: Logs error messages. ■ WARN: Logs error and warning messages. ■ INFO: Logs error, warning, and informational messages. ■ TRACE: Logs error, warning, informational, and trace messages. Trace logging is verbose. Use for initial configuration or troubleshooting. Type and dimension: string-scalar Default: INFO Example: "TRACE"

Table 5-1 Required attributes

Required attribute	Description
PidFile	This attribute is required when you want to enable the detection of a graceful shutdown outside of VCS control. See "PidFile" on page 147.
EnvFile	This attribute may be required when you use IBM HTTP Server. See "EnvFile" on page 147.

Table 5-2 Optional attributes

Optional attribute	Description
DirectiveAfter	A list of directives that httpd processes after reading the configuration file. Type and dimension: string-association Example: DirectiveAfter{} = { KeepAlive=On }
DirectiveBefore	A list of directives that httpd processes before it reads the configuration file. Type and dimension: string-association Example: DirectiveBefore{} = { User=nobody, Group=nobody }
User	Account name the agent uses to execute the httpd program. If you do not specify this value, the agent executes httpd as the root user. Type and dimension: string-scalar Example: "apache1"

Table 5-2 Optional attributes

Optional attribute	Description
EnableSSL	Set to 1 (true) to have the online agent function add support for SSL by including the option `-DSSL` in the start command. For example: `/usr/sbin/httpd -f path_to_httpd.conf -k start -DSSL` Where path_to_httpd.conf file is the path to the httpd.conf file. Set to 0 (false) it excludes the `-DSSL` option from the command. Type and dimension: boolean-scalar Default: 0 Example: "1"
HostName	The virtual host name that is assigned to the Apache server instance. The host name is used in second-level monitoring to establish a socket connection with the Apache HTTP server. **Note:** The HostName attribute is only required when the value of SecondLevelMonitor is 1 (true). Type and dimension: string-scalar Example: "web1.symantec.com"
Port	Port number where the Apache HTTP server instance listens. The port number is used in second-level monitoring to establish a socket connection with the server. Specify this attribute only if SecondLevelMonitor is set to 1 (true). Type and dimension: integer-scalar Default: 80 Example: "80"

Table 5-2 Optional attributes

Optional attribute	Description
EnvFile	Full path and file name of the file that is sourced before executing httpdDir/httpd. With Apache 2.0, the file *ServerRoot*/bin/envvars, which is supplied in most Apache 2.0 distributions, is commonly used to set the environment before executing httpd. Specifying this attribute is optional. If EnvFile is specified, the shell for user root must be Bourne, Korn, or C shell. This attribute may be required when you use the IBM HTTP Server if the online action fails. For example: set the EnvFile to /usr/IBM/HTTPServer/bin/envvars. Type and dimension: string-scalar Example: "/apache/server1/bin/envvars"
PidFile	The PidFile attribute sets the file to which the server records the process ID of the daemon. The value of PidFile attribute must be the absolute path where the Apache instance records the pid. This attribute is required when you want the agent to detect the graceful shutdown of the Web server. For the agent to detect the graceful shutdown of the Web server, the value of the IntentionalOffline resource type attribute must be 1 (true). Type and dimension: string-scalar Example: /var/run/httpd.pid
SharedObjDir	Full path of the directory in which the Apache HTTP shared object files are located. Specifying this attribute is optional. It is used when the HTTP Server is compiled using the SHARED_CORE rule. If you specify this attribute, the directory is passed to the `-R` option when executing the httpd program. Refer to the httpd man pages for more information about the `-R` option. Type and dimension: boolean-scalar Example: "/apache/server1/libexec"

Table 5-2 Optional attributes

Optional attribute	Description
SecondLevelTimeout	The number of seconds that the monitor agent function waits on the execution of second-level monitor. If the second-level monitor program does not return to calling the monitor agent function before the SecondLevelTimeout window expires, the monitor agent function no longer blocks on the program sub-process. It does, however, report that the resource is offline. The value should be high enough to allow the second level monitor enough time to complete. The value should be less than the value of the agent's MonitorTimeout. Type and dimension: integer-scalar Default: 30

Resource type definition

```
type Apache (
    static keylist SupportedActions = { "checkconffile.vfd" }
    static str ArgList[] = { ResLogLevel, State, IState, httpdDir,
    SharedObjDir, EnvFile, PidFile, HostName, Port, User,
    SecondLevelMonitor, SecondLevelTimeout, ConfigFile, EnableSSL,
    DirectiveAfter, DirectiveBefore }
    str ResLogLevel = INFO
    str httpdDir
    str SharedObjDir
    str EnvFile
    str PidFile
    str HostName
    int Port = 80
    str User
    boolean SecondLevelMonitor
    int SecondLevelTimeout = 30
    str ConfigFile
    boolean EnableSSL
    str DirectiveAfter{}
    str DirectiveBefore{}
    static int IntentionalOffline = 0
)
```

Apache Web server notes

The Apache Web server has the following notes:

Tasks to perform before you use the Apache Web server agent

Before you use this agent, perform the following tasks:

- Install the Apache server on shared or local disks.
- Ensure that you are able to start the Apache Web server outside of VCS control, with the specified parameters in the Apache configuration file (for example: /etc/apache/httpd.conf). For more information on how to start the server:
  See "About bringing an Apache Web server online outside of VCS control" on page 150.
- Specify the location of the error log file in the Apache configuration file for your convenience (for example: ErrorLog /var/apache/logs/error_log).
- Verify that the floating IP has the same subnet as the cluster systems.
- If you use a port other than the default 80, assign an exclusive port for the Apache server.
- Verify that the Apache server configuration files are identical on all cluster systems.
- Verify that the Apache server does not autostart on system startup.
- Verify that `Inetd` does not invoke the Apache server.
- Install the ACC Library 4.1.04.0 (VRTSacclib) if it is not already installed. If the ACC Library needs to be installed or updated, the library and its documentation can be obtained from the agent software media.

- Remove previous versions of this agent.
- The service group has disk and network resources to support the Apache server resource.
- Assign virtual host name and port to Apache Server.

Detecting application failure

The agent provides two methods to evaluate the state of an Apache HTTP server instance. The first state check is mandatory and the second is optional.

The first check determines the state of the Apache HTTP server. The check determines the state by searching for the existence of the parent httpd daemon. It also searches for at least one child httpd daemon. If the parent process and at least one child do not exist, VCS reports the resource as offline. If they do exist, and if the agent attribute SecondLevelMonitor is set to true, then a socket connection is established with the Apache HTTP server using the values specified by the Host and Port agent attributes. When connected, the agent issues an HTTP request to the server to test its ability to respond. If the HTTP Server responds with a return code between 0 and 408, the agent considers the server online. If the server fails to respond or returns any other code, the agent considers the server offline.

About bringing an Apache Web server online outside of VCS control

When you bring an Apache Web server online outside of VCS control, first source its environment file. Start the server with the `-f` option so the server knows which instance to start. You can then specify additional options (such as EnableSSL or SharedObjDir) that you want the server to use at start.

To start an Apache Web server outside of VCS control

1. Source the environment file if required.
2. Start the Apache Web server. You must use the `-f` option so that the agent can distinguish different instances of the server.
   httpdDir`/httpd -f` *ConfigFile* `-k start`
   Where *httpdDir* is /apache/v2.2/bin *ConfigFile* is /apache/v2.2/conf/httpd.conf. When fully formed, the start example looks like:
   `/apache/v2.2/bin/httpd -f /apache/v2.2/conf/httpd.conf -k start`
3. Specify additional options such as EnableSSL or SharedObjDir that you want to use when you start server. When you add EnableSSL to the command, it resembles:
   httpdDir`/httpd -f` *ConfigFile* `-k start -DSSL`

About the ACC Library

The agent functions for the Apache HTTP server depend on a set of Perl modules that are known as the ACC Library. The ACC Library contains the common, reusable functions that perform tasks such as process identification, logging, and system calls.

When you install the ACC library in a VCS environment, you must install the ACC library package before you install the agent.

To install or update the ACC library package, locate the library and related documentation on the agent disc and in the compressed agent tar file.

High Availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node. These discrepancies might prevent a service group from going online on a specific node. For Apache resources, when the Apache Web server is installed locally, the high availability fire drill checks for the validity of these attributes:

- ConfigFile
- httpdDir

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Sample configurations

```
group ApacheG1(
        SystemList = { host1 = 0, host2 = 1 }
        )

        Apache httpd_server (
                Critical = 0
                httpdDir = "/apache/bin"
                HostName = vcsaix1
                Port = 8888
                User = root
                SecondLevelMonitor = 1
                ConfigFile = "/apache/conf/httpd.conf"
                )

        DiskGroup Apache_dg (
                Critical = 0
                DiskGroup = apc1
                )
```

```
IP Apache_ip (
        Critical = 0
        Device = en0
        Address = "11.123.99.168"
        NetMask = "255.255.254.0"
        )

Mount Apache_mnt (
        Critical = 0
        MountPoint = "/apache"
        BlockDevice = "/dev/vx/dsk/apc1/apcvol1"
        FSType = vxfs
        FsckOpt = "-y"
        )

Apache_mnt requires Apache_dg
httpd_server requires Apache_mnt
httpd_server requires Apache_ip
```

Application agent

The Application agent brings 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. The executables must exist locally on each node. You can use this agent to provide high availability for applications that do not have custom agents.

An application runs in the default context of root. Specify the user name to run an application in a user context.

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

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node. These discrepancies might prevent a service group from going online on a specific node. For Application resources, the high availability fire drill checks for:

- The availability of the specified program
- Execution permissions for the specified program
- The existence of the specified user on the host
- The existence of the same binary on all nodes

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Dependencies

Depending on how you plan to use it, this type of resource can depend on IP, IPMultiNIC, and Mount resources.

Figure 5-2 Sample service group for an Application resource

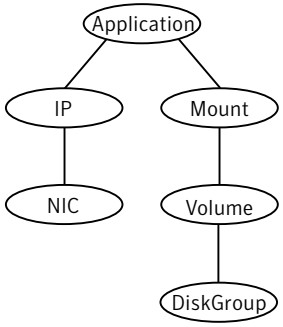

Agent functions

Online	Runs the StartProgram attribute with the specified parameters in the context of the specified user.
Offline	Runs the StopProgram attribute with the specified parameters in the context of the specified user.
Monitor	If you specify the MonitorProgram attribute, the agent executes the user-defined MonitorProgram in the user-specified context. If you specify the PidFiles attribute, the routine verifies that the process ID that is found in each listed file is running. If you specify the MonitorProcesses attribute, the routine verifies that each listed process is running in the context you specify. Use any combination among these attributes (MonitorProgram, PidFiles, or MonitorProcesses) to monitor the application. If any one the processes that are specified in either PidFiles or MonitorProcesses is determined not to be running, the monitor returns OFFLINE. If the process terminates ungracefully, the monitor returns OFFLINE and failover occurs.
Clean	Terminates processes specified in PidFiles or MonitorProcesses. Ensures that only those processes (that are specified in the MonitorProcesses attribute) running with the user ID specified in the User attribute are killed. If the CleanProgram is defined, the agent executes the CleanProgram.

State definitions

ONLINE	Indicates that all processes that are specified in the PidFiles and the MonitorProcesses attribute are running and that the MonitorProgram returns ONLINE.
OFFLINE	Indicates that at least one process that are specified in the PidFiles attribute or MonitorProcesses is not running, or that the MonitorProgram returns OFFLINE.
UNKNOWN	Indicates an indeterminable application state or invalid configuration.

Attributes

Table 5-3 Required attributes

Required attribute	Description
StartProgram	The executable, created locally on each node, which starts the application. Specify the complete path of the executable. Applicable command line arguments follow the name of the executable and have spaces separating them. **Note:** Do not use the opening and closing ({ }) brace symbols in this string. Type and dimension: string-scalar Example: "/usr/sbin/samba start"
StopProgram	The executable, created locally on each node, which stops the application. Specify the complete path of the executable. Applicable command line arguments follow the name of the executable and have spaces separating them. **Note:** Do not use the opening and closing ({ }) brace symbols in this string. Type and dimension: string-scalar Example: "/usr/sbin/samba stop"
At least one of the following attributes: ■ MonitorProcesses ■ MonitorProgram ■ PidFiles	See “Optional attributes” on page 157.

Table 5-4 Optional attributes

Optional attribute	Description
CleanProgram	The executable, created locally on each node, which forcibly stops the application. Specify the complete path of the executable. Applicable command line arguments follow the name of the executable and have spaces separating them. Type and dimension: string-scalar
MonitorProcesses	A list of processes that you want monitored and cleaned. Each process name is the name of an executable. Qualify the executable name with its complete path if the path starts the executable. The process name must be the full command line argument that the `ps -u user -eo pid,comm | more` command displays for the process. Type and dimension: string-vector
MonitorProgram	The executable, created locally on each node, which monitors the application. Specify the complete path of the executable. Applicable command line arguments follow the name of the executable and have spaces separating them. MonitorProgram can return the following VCSAgResState values: OFFLINE value is 100; ONLINE values range from 101 to 110 (depending on the confidence level); 110 equals confidence level of 100%. Any other value = UNKNOWN. **Note:** Do not use the opening and closing ({ }) brace symbols in this string. Type and dimension: string-scalar Example: "/usr/sbin/sample_app_monitor all"

Table 5-4 Optional attributes

Optional attribute	Description
PidFiles	A list of PID (process ID) files that contain the PID of the processes that you want monitored and cleaned. These are application generated files. Each PID file contains one monitored PID. Specify the complete path of each PID file in the list. The process ID can change when the process restarts. If the application takes time to update the PID file, the agent's Monitor function may return an incorrect result. If incorrect results occur, increase the ToleranceLimit in the resource definition. Type and dimension: string-vector
User	The user ID for running StartProgram, StopProgram, MonitorProgram, and CleanProgram. The processes that are specified in the MonitorProcesses list must run in the context of the specified user. Monitor checks the processes to make sure they run in this context. Type and dimension: string-scalar Default: root

Resource type definition

```
type Application (
    static keylist SupportedActions = { "program.vfd", "user.vfd",
    "cksum.vfd", getcksum }
    static str ArgList[] = { User, StartProgram, StopProgram,
    CleanProgram, MonitorProgram, PidFiles, MonitorProcesses }
    str User
    str StartProgram
    str StopProgram
    str CleanProgram
    str MonitorProgram
    str PidFiles[]
    str MonitorProcesses[]
)
```

Sample configurations

Configuration 1

In this example, you configure the executable samba as StartProgram and StopProgram, with start and stop specified as command line arguments respectively. Configure the agent to monitor two processes: a process that the smbd.pid specifies and the process nmbd.

```
Application samba_app (
    User = "root"
    StartProgram = "/usr/sbin/samba start"
    StopProgram = "/usr/sbin/samba stop"
    PidFiles = { "/var/lock/samba/smbd.pid" }
    MonitorProcesses = { "nmbd" }
)
```

Configuration 2

In this example, since no user is specified, it uses the root user. The executable samba starts and stops the application using start and stop as the command line arguments. The executable sambaMonitor monitors the application and uses `all` as its command line argument. The agent also monitors the smbd and nmbd processes.

```
Application samba_app2 (
    StartProgram = "/usr/sbin/samba start"
    StopProgram = "/usr/sbin/samba stop"
    CleanProgram = "/usr/sbin/samba force stop"
    MonitorProgram = "/usr/local/bin/sambaMonitor all"
    MonitorProcesses = { "smbd", "nmbd" }
)
```

Process agent

The Process agent starts, stops, and monitors a process that you specify. You can use the agent to make a process highly available or to monitor it.

High availability fire drill

The high availability fire drill detects discrepancies between the VCS configuration and the underlying infrastructure on a node; discrepancies that might prevent a service group from going online on a specific node. For Process resources, the high availability fire drill checks for:

- The existence of the specified process
- Execution permissions for the specified process
- The existence of a binary executable for the specified process
- The existence of the same binary on all nodes

For more information about using the high availability fire drill see the *Veritas Cluster Server User's Guide*.

Dependencies

Depending on the context, this type of resource can depend on IP, IPMultiNIC, and Mount resources.

Figure 5-3 Sample service group for a Process resource

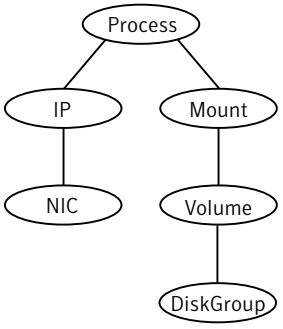

Agent functions

Online	Starts the process with optional arguments.
Offline	Terminates the process with a `SIGTERM`. If the process does not exit, a `SIGKILL` is sent.
Monitor	Checks to see if the process is running by scanning the process table for the name of the executable pathname and argument list.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	Indicates that the specified process is running in the specified user context.The agent only reports the process as online if the value configured for PathName attribute exactly matches the process listing from the ps output.
OFFLINE	Indicates that the specified process is not running in the specified user context.
FAULTED	Indicates that the process has terminated unexpectedly.
UNKNOWN	Indicates that the agent can not determine the state of the process.

Attributes

Table 5-5 Required attribute

Required attribute	Description
PathName	Complete pathname to access an executable program. This path includes the program name. If a script controls the process, the PathName defines the complete path to the shell. Type and dimension: string-scalar Example: "/usr/sbin/sendmail"

Table 5-6 Optional attributes

Optional attribute	Description
Arguments	Passes arguments to the process. If a script controls the process, the script is passed as an argument. Separate multiple arguments with a single space. A string cannot accommodate more than one space between arguments, nor allow for leading or trailing whitespace characters. This attribute must not exceed 80 characters. Type and dimension: string-scalar Example: "bd q1h"

Resource type definition

```
type Process (
    static keylist SupportedActions = { "program.vfd", getcksum }
    static str ArgList[] = { PathName, Arguments }
    str PathName
    str Arguments
)
```

Sample configurations

Configuration 1

```
Process  usr_lib_sendmail (
    PathName = "/usr/lib/sendmail"
    Arguments = "bd q1h"
    )
```

Configuration 2

```
include "types.cf"
cluster ProcessCluster (
.
.
.
group ProcessGroup (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )

    Process Process1 (
        PathName = "/usr/local/bin/myprog"
        Arguments = "arg1 arg2"
        )

    Process Process2 (
    PathName = "/bin/csh"
    Arguments = "/tmp/funscript/myscript"
    )

    // resource dependency tree
    //
    //    group ProcessGroup
    //    {
    //    Process Process1
    //    Process Process2
    //    }
```

ProcessOnOnly agent

The ProcessOnOnly agent starts and monitors a process that you specify. You can use the agent to make a process highly available or to monitor it. This resource's Operation value is OnOnly.

VCS uses this agent internally to mount security processes in a secure cluster.

Dependencies

No child dependencies exist for this resource.

Agent functions

Online	Starts the process with optional arguments.
Monitor	Checks to see if the process is alive by scanning the process table for the name of the executable pathname and argument list.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	Indicates that the specified process is running.The agent only reports the process as ONLINE if the value configured for PathName attribute exactly matches the process listing from the ps output.
FAULTED	Indicates that the process has unexpectedly terminated.
UNKNOWN	Indicates that the agent can not determine the state of the process.

Attributes

Table 5-7 Required attributes

Required attribute	Description
PathName	Defines complete pathname to access an executable program. This path includes the program name. If a process is controlled by a script, the PathName defines the complete path to the shell. The value configured for this attribute needs to match the process listing from the ps output for the agent to display as ONLINE. Type and dimension: string-scalar

Table 5-8 Optional attributes

Optional attribute	Description
Arguments	Passes arguments to the process. If a process is controlled by a script, the script is passed as an argument. Multiple arguments must be separated by a single space. A string cannot accommodate more than one space between arguments, nor allow for leading or trailing whitespace characters. Arguments must not exceed 80 characters (total). Type and dimension: string-scalar
IgnoreArgs	A flag that indicates whether monitor ignores the argument list. ■ If the value is 0, it checks the process pathname and argument list. ■ If the value is 1, it only checks for the executable pathname and ignores the rest of the argument list. Type and dimension: boolean-scalar Default: 0

Resource type definition

```
type ProcessOnOnly (
    static str ArgList[] = { IgnoreArgs, PathName, Arguments }
    static str Operations = OnOnly
    int IgnoreArgs
    str PathName
    str Arguments
)
```

Sample configurations

```
group VxSS (
    SystemList = { north = 0, south = 1 }
    Parallel = 1
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

Phantom phantom_vxss (
    )

ProcessOnOnly vxatd (
    IgnoreArgs = 1
    PathName = "/opt/VRTSat/bin/vxatd"
    )

// resource dependency tree
//
// group VxSS
```

Chapter 6

Infrastructure and support agents

This chapter contains the following agents:

About the infrastructure and support agents

Use the infrastructure and support agents to monitor Veritas components and VCS objects.

NotifierMngr agent

Starts, stops, and monitors a notifier process, making it highly available. The notifier process manages the reception of messages from VCS and the delivery of those messages to SNMP consoles and SMTP servers. See the *Veritas Cluster Server User's Guide* for a description of types of events that generate notification. See the `notifier`(1) manual page to configure notification from the command line.

You cannot dynamically change the attributes of the NotifierMngr agent using the `hares -modify` command. Changes made using this command are only effective after restarting the notifier.

Dependency

The NotifierMngr resource can depend on the NIC resource.

Agent functions

Online	Starts the notifier process with its required arguments.
Offline	VCS sends a `SIGABORT`. If the process does not exit within one second, VCS sends a `SIGKILL`.
Monitor	Monitors the notifier process.
Clean	Sends `SIGKILL`.

State definitions

ONLINE	Indicates that the Notifier process is running.
OFFLINE	Indicates that the Notifier process is not running.
UNKNOWN	Indicates that the user did not specify the required attribute for the resource.

Attributes

Table 6-1 Required attributes

Required attribute	Description
SnmpConsoles	Specifies the machine names of the SNMP managers and the severity level of the messages to be delivered. The severity levels of messages are Information, Warning, Error, and SevereError. Specifying a given severity level for messages generates delivery of all messages of equal or higher severity. **Note:** SnmpConsoles is a required attribute if SmtpServer is not specified; otherwise, SnmpConsoles is an optional attribute. Specify both SnmpConsoles and SmtpServer if desired. Type and dimension: string-association Example: "172.29.10.89" = Error, "172.29.10.56" = Information
SmtpServer	Specifies the machine name of the SMTP server. **Note:** SmtpServer is a required attribute if SnmpConsoles is not specified; otherwise, SmtpServer is an optional attribute. You can specify both SmtpServer and SnmpConsoles if desired. Type and dimension: string-scalar Example: "smtp.your_company.com"

Table 6-2 Optional attributes

Optional attribute	Description
MessagesQueue	Size of the VCS engine's message queue. Minimum value is 30. Type and dimension: integer-scalar Default: 30

Table 6-2 Optional attributes

Optional attribute	Description
NotifierListeningPort	Any valid, unused TCP/IP port number. Type and dimension: integer-scalar Default: 14144
SmtpFromPath	Set to a valid email address, if you want the notifier to use a custom email address in the FROM: field. Type and dimension: string-scalar Example: "usera@example.com"
SmtpRecipients	Specifies the email address where SMTP sends information and the severity level of the messages. The severity levels of messages are Information, Warning, Error, and SevereError. Specifying a given severity level for messages indicates that all messages of equal or higher severity are received. **Note:** SmtpRecipients is a required attribute if you specify SmtpServer. Type and dimension: string-association Example: "james@symantec.com" = SevereError, "admin@symantec.com" = Warning
SmtpReturnPath	Set to a valid email address, if you want the notifier to use a custom email address in the Return-Path: <> field. If the mail server specified in SmtpServer does not support VRFY, then you need to set the SmtpVrfyOff to `1` in order for the SmtpReturnPath value to take effect. Type and dimension: string-scalar Example: "usera@example.com"

Table 6-2 Optional attributes

Optional attribute	Description
SmtpServerTimeout	This attribute represents the time in seconds notifier waits for a response from the mail server for the SMTP commands it has sent to the mail server. This value can be increased if you notice that the mail server is taking a longer duration to reply back to the SMTP commands sent by notifier. Type and dimension: integer-scalar Default: 10
SmtpServerVrfyOff	Set this value to `1` if your mail server does not support SMTP VRFY command. If you set this value to `1`, the notifier does not send a SMTP VRFY request to the mail server specified in SmtpServer attribute while sending emails. Type and dimension: boolean-scalar Default: 0
SnmpCommunity	Specifies the community ID for the SNMP manager. Type and dimension: string-scalar Default: public
SnmpdTrapPort	Port on the SNMP console machine where SNMP traps are sent. If you specify more than one SNMP console, all consoles use this value. Type and dimension: integer-scalar Default: 162
EngineListeningPort	Change this attribute if the VCS engine is listening on a port other than its default port. Type and dimension: integer-scalar Default: 14141

Resource type definition

```
type NotifierMngr (
    static int RestartLimit = 3
    static str ArgList[] = { EngineListeningPort, MessagesQueue,
    NotifierListeningPort, SnmpdTrapPort, SnmpCommunity,
    SnmpConsoles, SmtpServer, SmtpServerVrfyOff,
    SmtpServerTimeout, SmtpReturnPath, SmtpFromPath, SmtpRecipients
    }
    int EngineListeningPort = 14141
    int MessagesQueue = 30
    int NotifierListeningPort = 14144
    int SnmpdTrapPort = 162
    str SnmpCommunity = "public"
    str SnmpConsoles{}
    str SmtpServer
    boolean SmtpServerVrfyOff = 0
    int SmtpServerTimeout = 10
    str SmtpReturnPath
    str SmtpFromPath
    str SmtpRecipients{}
)
```

Sample configuration

In the following configuration, the NotifierMngr agent is configured to run with two resource groups: NicGrp and Grp1. NicGrp contains the NIC resource and a Phantom resource that enables VCS to determine the online and offline status of the group. See the Phantom agent for more information on verifying the status of groups that only contain OnOnly or Persistent resources such as the NIC resource. You must enable NicGrp to run as a parallel group on both systems.

Grp1 contains the NotifierMngr resource (ntfr) and a Proxy resource (nicproxy), configured for the NIC resource in the first group.

In this example, NotifierMngr has a dependency on the Proxy resource.

Note: Only one instance of the notifier process can run in a cluster. The process cannot run in a parallel group.

The NotifierMngr resource sets up notification for all events to the SNMP console `snmpserv`. In this example, only messages of SevereError level are sent to the SMTP server (`smtp.example.com`), and the recipient (`vcsadmin@example.com`).

Configuration

```
system north

system south

group  NicGrp (
    SystemList = { north, south }
    AutoStartList = { north }
    Parallel = 1
    )

    Phantom my_phantom (
    )

    NIC    NicGrp_en0 (
        Enabled = 1
        Device  = en0
        NetworkType = ether
        )

group Grp1 (
    SystemList = { north, south }
    AutoStartList = { north }
    )
```

```
Proxy nicproxy(
TargetResName = "NicGrp_en0"
)

NotifierMngr ntfr (
    SnmpConsoles = { snmpserv = Information }
    SmtpServer = "smtp.example.com"
    SmtpRecipients = { "vcsadmin@example.com" =
    SevereError }
)

ntfr requires nicproxy


// resource dependency tree
//
//      group Grp1
//      {
//      NotifierMngr ntfr
//              {
//              Proxy nicproxy
//              }
//      }
```

VRTSWebApp agent

Brings Web applications online, takes them offline, and monitors their status. This agent is used to monitor the Web consoles of various Symantec products, such as the Cluster Management Console.

Agent functions

Online	Starts the Web application with the specified parameters. If the Web server is not already running, it first starts the server.
Offline	Removes the Web application from the Web server. If no other Web application is running, it shuts down the Web server.
Monitor	Checks if the specified Web application is currently running inside the Web server. If the application is running, monitor reports ONLINE. If the application is not running, monitor reports OFFLINE.
Clean	Removes the Web application from the Web server. If no other Web application is running, it shuts down the Web server.

State definitions

ONLINE	Indicates that the Web application is running.
OFFLINE	Indicates that the Web application is not running.
UNKNOWN	Indicates that the agent could not determine the state of the resource or that the resource attributes are invalid.

Attributes

Table 6-3 Required attributes

Required attribute	Description
AppName	Name of the application as it appears in the Web server. Type and dimension: string-scalar Example: "cmc"
InstallDir	Path to the Web application installation. You must install the Web application as a .war file with the same name as the AppName parameter. Point this attribute to the directory that contains this .war file. Type and dimension: string-scalar Example: If the AppName is cmc and InstallDir is /opt/VRTSweb/VERITAS, the agent constructs the path for the Web application as /opt/VRTSweb/VERITAS/cmc.war.
TimeForOnline	The time the Web application takes to start after loading it into the Web server. This parameter is returned as the exit value of the online script, which inform VCS of the time it needs to wait before calling monitor on the Web application resource. This attribute is typically at least five seconds. Type and dimension: integer-scalar

Table 6-4 Optional attributes

Optional attribute	Description
NumThreads	Number of threads used within the agent process for managing resources. This number does not include threads used for other internal purposes. Symantec strongly recommends that you retain the default value of the NumThreads attribute of 1. Setting this attribute to a higher value may result in agent function timeouts due to serialization of underlying commands. Default: 1

Resource type definition

```
type VRTSWebApp (
    static str ArgList[] = { AppName, InstallDir, TimeForOnline }
    str AppName
    str InstallDir
    int TimeForOnline
    static int NumThreads = 1
)
```

Sample configuration

```
VRTSWebApp VCSweb (
    AppName = "cmc"
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
)
```

Proxy agent

The Proxy agent mirrors the state of another resource on a local or remote system. It provides a means to specify and modify one resource and have its state reflected by its proxies. You can use the agent when you need to replicate the status of a resource.

A Proxy resource can only point to None or OnOnly type of resources, and can reside in a failover/parallel group. A target resource and its proxy cannot be in the same group.

Dependencies

No dependencies exist for the Proxy resource.

Figure 6-1 Sample service group for an Proxy resource

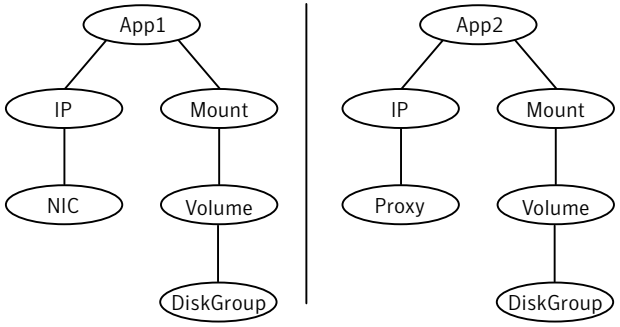

Agent functions

Monitor Determines status based on the target resource status.

Attributes

Table 6-5 Required attribute

Required attribute	Description
TargetResName	Name of the target resource that the Proxy resource mirrors. The target resource must be in a different resource group than the Proxy resource. Type and dimension: string-scalar Example: "tmp_VRTSvcs_file1"

Table 6-6 Optional attribute

Optional attribute	Description
TargetSysName	Mirrors the status of the TargetResName attribute on systems that the TargetSysName variable specifies. If this attribute is not specified, the Proxy resource assumes the system is local. Type and dimension: string-scalar Example: "sysa"

Resource type definition

```
type Proxy (
    static str ArgList[] = { TargetResName, TargetSysName,
    "TargetResName:Probed", "TargetResName:State" }
    static int OfflineMonitorInterval = 60
    static str Operations = None
    str TargetResName
    str TargetSysName
)
```

Sample configurations

Configuration 1

The proxy resource mirrors the state of the resource tmp_VRTSvcs_file1 on the local system.

```
Proxy proxy1 (
    TargetResName = "tmp_VRTSvcs_file1"
)
```

Configuration 2

The proxy resource mirrors the state of the resource tmp_VRTSvcs_file1 on sysa.

```
Proxy proxy1(
    TargetResName = "tmp_VRTSvcs_file1"
    TargetSysName = "sysa"
)
```

Configuration

The proxy resource mirrors the state of the resource mnic on the local system; note that target resource is in grp1, and the proxy is in grp2; a target resource and its proxy cannot be in the same group.

```
group grp1 (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    )

MultiNICA mnic (
        Device@sysa = { en0 = "10.128.8.42", en1 = "10.128.8.42" }
        Device@sysb = { en0 = "10.128.8.43", en1 = "10.128.8.43" }
        NetMask = "255.255.255.0"
        NameServerAddr = "10.130.8.1"
        Gateway = "10.128.1.1"
        Domain = "veritas.com"
        BroadcastAddr = "10.128.25.255"
```

```
        Options = "mtu m"
        )

        IPMultiNIC ip1 (
            Address = "166.98.14.78"
            NetMask = "255.255.255.0"
            MultiNICAResName = mnic
            Options = "mtu m"
            )
    ip1 requires mnic

group grp2 (
        SystemList = { sysa, sysb }
        AutoStartList = { sysa }
        )
        IPMultiNIC ip2 (
            Address = "166.98.14.79"
            NetMask = "255.255.255.0"
            MultiNICAResName = mnic
            Options = "mtu m"
            )
        Proxy proxy (
            TargetResName = mnic
            )
    ip2 requires proxy
```

Phantom agent

The Phantom agent enables VCS to determine the status of parallel service groups that do not include OnOff resources. Do not use the Phantom agent in failover service groups. You can use the agent to determines the state of service groups having resources of type None only.

Do not attempt manual online or offline operations on the Phantom resource or on the service group containing the Phantom resource. Doing so may result in unpredictable behavior.

Dependencies

No dependencies exist for the Phantom resource.

Figure 6-2 Sample service group for a Phantom resource

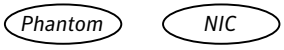

Agent functions

Monitor Determines status based on the status of the service group.

Attribute

Table 6-7 Attribute

Attribute	Description
Dummy	The Dummy attribute is for internal use only.

Resource type definition

```
type Phantom (
    static str ArgList[] = { Dummy }
    str Dummy
)
```

Sample configurations

Configuration 1

```
Phantom boo (
)
```

Configuration 2

The following example shows a complete main.cf, in which the FileNone resource and the Phantom resource are in the same group.

```
include "types.cf"

cluster PhantomCluster

system sysa

system sysb

group phantomgroup (
    SystemList = { sysa, sysb }
    AutoStartList = { sysa }
    Parallel = 1
    )

    FileNone my_file_none (
        PathName = "/tmp/file_none"
        )
    Phantom my_phantom (
        )

    // resource dependency tree
    //
    //    group maingroup
    //    {
    //    Phantom my_Phantom
    //    FileNone my_file_none
    //    }
```

RemoteGroup agent

The RemoteGroup agent establishes dependencies between applications that are configured on different VCS clusters. For example, you configure an Apache resource in a local cluster, and a MySQL resource in a remote cluster. In this example, the Apache resource depends on the MySQL resource. You can use the RemoteGroup agent to establish this dependency between these two resources.

With the RemoteGroup agent, you can monitor or manage a service group that exists in a remote cluster. Some points about configuring the RemoteGroup resource follow:

- For each remote service group that you want to monitor or manage, you must configure a corresponding RemoteGroup resource in the local cluster.
- Multiple RemoteGroup resources in a local cluster can manage corresponding multiple remote service groups in different remote clusters.
- You can include the RemoteGroup resource in any kind of resource or service group dependency tree.
- A combination of the state of the local service group and the state of the remote service group determines the state of the RemoteGroup resource.

Symantec supports the RemoteGroup agent when it points to a global group. The RemoteGroup agent must then map the state of the global group in the local cluster.

For more information on the functionality of this agent see the *Veritas Cluster Server User's Guide*.

Dependency

As a best practice, establish a RemoteGroup resource dependency on a NIC resource. Symantec recommends that the RemoteGroup resource not be by itself in a service group.

Agent functions

Online	Brings the remote service group online. See the "ControlMode" on page 187 for more information.
Offline	Takes the remote service group offline. See the "ControlMode" on page 187 for more information.
Monitor	Monitors the state of the remote service group. The true state of the remote service group is monitored only on the online node in the local cluster. See the "VCSSysName" on page 186.
Clean	If the RemoteGroup resource faults, the Clean function takes the remote service group offline. See the "ControlMode" on page 187 for more information.

State definitions

ONLINE	Indicates that the remote service group is either in an ONLINE or PARTIAL state.
OFFLINE	Indicates that the remote service group is in an OFFLINE or FAULTED state. The true state of the remote service group is monitored only on the online node in the local cluster.
FAULTED	Indicates that the RemoteGroup resource has unexpectedly gone offline.
UNKNOWN	Indicates that a problem exists either with the configuration or the ability of the RemoteGroup resource to determine the state of the remote service group.

Attributes

Table 6-8 Required attributes

Required attribute	Description
IpAddress	The IP address or DNS name of a node in the remote cluster. The IP address can be either physical or virtual. When configuring a virtual IP address of a remote cluster, do not configure the IP resource as a part of the remote service group. Type and dimension: string-scalar Examples: "www.example.com" or "11.183.12.214"
Port	This is a required attribute when the remote cluster listens on a port other than the default value of 14141. See "Port" on page 189.
GroupName	The name of the service group on the remote cluster that you want the RemoteGroup agent to monitor or manage. Type and dimension: string-scalar Example: "DBGrp"
VCSSysName	You must set this attribute to either the VCS system name or the ANY value. ■ ANY The RemoteGroup resource goes online if the remote service group is online on any node in the remote cluster. ■ *VCSSysName* Use the name of a VCS system in a remote cluster where you want the remote service group to be online when the RemoteGroup resource goes online. Use this to establish a one-to-one mapping between the nodes of the local and remote clusters. Type and dimension: string-scalar Example: "vcssys1" or "ANY"

Table 6-8 Required attributes

Required attribute	Description
ControlMode	Select only one of these values to determine the mode of operation of the RemoteGroup resource: MonitorOnly, OnlineOnly, or OnOff. ■ OnOff The RemoteGroup resource brings the remote service group online or takes it offline. When you set the VCSSysName attribute to ANY, the SysList attribute of the remote service group determines the node where the remote service group onlines. ■ MonitorOnly The RemoteGroup resource only monitors the state of the remote service group. The RemoteGroup resource cannot online or offline the remote service group. Make sure that you bring the remote service group online before you online the RemoteGroup resource. ■ OnlineOnly The RemoteGroup resource only brings the remote service group online. The RemoteGroup resource cannot take the remote service group offline. When you set the VCSSysName attribute to ANY, the SysList attribute of the remote service group determines the node where the remote service group onlines. Type and dimension: string-scalar

Table 6-8 Required attributes

Required attribute	Description
Username	This is the login user name for the remote cluster. When you set the ControlMode attribute to OnOff or OnlineOnly, the Username must have administrative privileges for the remote service group that you specify in the GroupName attribute. When you use the RemoteGroup Wizard to enter your username data, you need to enter your username and the domain name in separate fields. For a cluster that has the Symantec Product Authentication Service, you do not need to enter the domain name. For a secure remote cluster: ■ Local Unix user user@nodename—where the nodename is the name of the node that is specified in the IpAddress attribute. Do not set the DomainType attribute. ■ NIS or NIS+ user user@domainName—where domainName is the name of the NIS or NIS+ domain for the user. You must set the value of the DomainType attribute to either to nis or nisplus. Type and dimension: string-scalar Example: ■ For a cluster without the Symantec Product Authentication Service: "johnsmith" ■ For a secure remote cluster: "foobar@example.com"
Password	This is the password that corresponds to the user that you specify in the Username attribute. You must encrypt the password with the `vcsencrypt -agent` command. **Note:** Do not use the vcsencrypt utility when entering passwords from a configuration wizard or from the Cluster Management Console or the Cluster Manager (Java Console). Type and dimension: string-scalar

Table 6-9 Optional attributes

Optional attribute	Description
DomainType	For a secure remote cluster only, enter the domain type information for the specified user. For users who have the domain type unixpwd, you do not have to set this attribute. Type: string-scalar Example: "nis", "nisplus"
BrokerIp	For a secure remote cluster only. If you need the RemoteGroup agent to communicate to a specific authentication broker, set the value of this attribute to the broker's IP address. Type: string-scalar Example: "128.11.295.51"
Port	The port where the remote engine listens for requests. This is an optional attribute, unless the remote cluster listens on a port other than the default value of 14141. Type and dimension: integer-scalar Default: 14141
OfflineWaitTime	The maximum expected time in seconds that the remote service group may take to offline. VCS calls the clean function for the RemoteGroup resource if the remote service group takes a longer time to offline than the time that you have specified for this attribute. Type and dimension: integer-scalar Default: 0

Table 6-10 Type-level attributes

Type level attributes	Description
OnlineRetryLimit OnlineWaitLimit ToleranceLimit MonitorInterval AutoFailover	In case of remote service groups that take a longer time to Online, Symantec recommends that you modify the default OnlineWaitLimit and OnlineRetryLimit attributes. If you expect the RemoteGroup agent to tolerate sudden offlines of the remote service group, then modify the ToleranceLimit attribute. See the *Veritas Cluster Server User's Guide* for more information about these attributes.

1 Resource type definition

```
type RemoteGroup (
    static int OnlineRetryLimit = 2
    static int ToleranceLimit = 1
    static str ArgList[] = { IpAddress, Port, Username, Password,
    GroupName, VCSSysName, ControlMode, OfflineWaitTime,
    DomainType, BrokerIp }
    str IpAddress
    int Port = 14141
    str Username
    str Password
    str GroupName
    str VCSSysName
    str ControlMode
    int OfflineWaitTime
    str DomainType
    str BrokerIp
)
```

Chapter 7

Testing agents

This chapter contains the following agents:

About the program support agents

Use the program support agents to provide high availability for program support resources.

ElifNone agent

The ElifNone agent monitors a file. It checks for the file's absence.

You can use the ElifNone agent to test service group behavior. You can also use it as an impostor resource, where it takes the place of a resource for testing.

Dependencies

No dependencies exist for the ElifNone resource.

Figure 7-1 Sample service group for an ElifNone resource

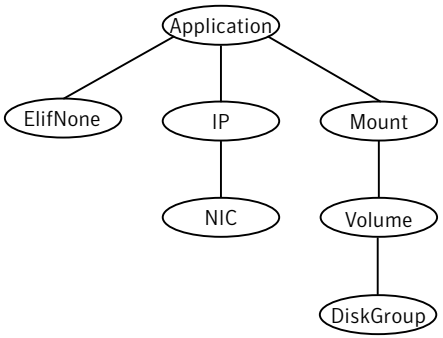

Agent function

Monitor — Checks for the specified file. If it exists, the resource faults. If it does not exist, the agent reports as ONLINE.

Attributes

Table 7-1 Required attribute

Required attribute	Description
PathName	Specifies the complete pathname. Starts with a slash (/) preceding the file name. Type and dimension: string-scalar Example: "/tmp/file01"

Resource type definition

```
type ElifNone (
    static str ArgList[] = { PathName }
    static int OfflineMonitorInterval = 60
    static str Operations = None
    str PathName
)
```

Sample configuration

```
ElifNone tmp_file01 (
    PathName = "/tmp/file01"
)
```

FileNone agent

Monitors a file—checks for the file's existence.

You can use the FileNone agent to test service group behavior. You can also use it as an "impostor" resource, where it takes the place of a resource for testing.

Dependencies

No dependencies exist for the FileNone resource.

Figure 7-2 Sample service group for an FileNone resource

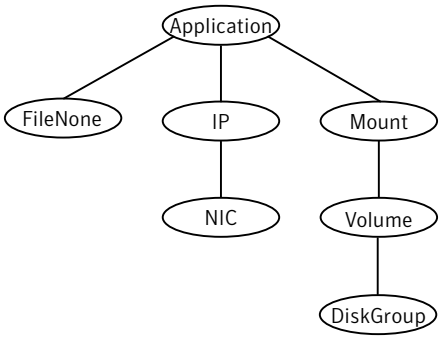

Agent functions

Monitor — Checks for the specified file. If it exists, the agent reports as ONLINE. If it does not exist, the resource faults.

Attribute

Table 7-2 Required attribute

Required attribute	Description
PathName	Specifies the complete pathname. Starts with a slash (/) preceding the file name. Type and dimension: string-scalar Example: "/tmp/file01"

Resource type definition

```
type FileNone (
    static int AutoRestart = 1
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { PathName }
    static str Operations = None
    str PathName
)
```

Sample configuration

```
FileNone tmp_file01 (
    PathName = "/tmp/file01"
)
```

FileOnOff agent

The FileOnOff agent creates, removes, and monitors files.

You can use the FileNone agent to test service group behavior. You can also use it as an “impostor” resource, where it takes the place of a resource for testing.

Dependencies

No dependencies exist for the FileOnOff resource.

Figure 7-3 Sample service group for a FileOnOff resource

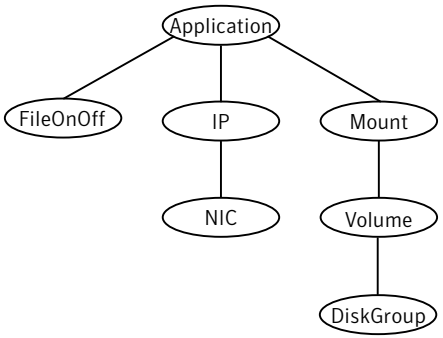

Agent functions

Online	Creates an empty file with the specified name if the file does not already exist.
Offline	Removes the specified file.
Monitor	Checks for the specified file. If it exists, the agent reports as ONLINE. If it does not exist, the agent reports as OFFLINE.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

Attribute

Table 7-3 Required attribute

Required attribute	Description
PathName	Specifies the complete pathname. Starts with a slash (/) preceding the file name. Type and dimension: string-scalar Example: "/tmp/file01"

Resource type definition

```
type FileOnOff (
    static str ArgList[] = { PathName }
    str PathName
)
```

Sample configuration

```
FileOnOff tmp_file01 (
    PathName = "/tmp/file01"
)
```

FileOnOnly agent

The FileOnOnly agent creates and monitors files.

You can use the FileNone agent to test service group behavior. You can also use it as an “impostor” resource, where it takes the place of a resource for testing.

Dependencies

No dependencies exist for the FileOnOnly resource.

Figure 7-4 Sample service group for a FileOnOnly resource

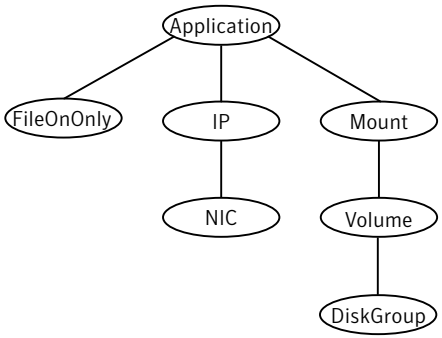

Agent functions

Online	Creates an empty file with the specified name, unless one already exists.
Monitor	Checks for the specified file. If it exists, the agent reports as ONLINE. If it does not exist, the resource faults.

Attribute

Table 7-4 Required attributes

Required attribute	Description
PathName	Specifies the complete pathname. Starts with a slash (/) preceding the file name. Type and dimension: string-scalar Example: "/tmp/file02"

Resource type definition

```
type FileOnOnly (
    static str ArgList[] = { PathName }
    static str Operations = OnOnly
    str PathName
)
```

Sample configuration

```
FileOnOnly tmp_file02 (
    PathName = "/tmp/file02"
)
```

Glossary

administrative IP address
The operating system controls these IP addresses and brings them up even before VCS brings applications online. Use them to access a specific system over the network for doing administrative tasks, for example: examining logs to troubleshoot issues, cleaning up temp files to free space, etc. Typically, you have one administrative IP address per node.

agent function
Agent functions start, stop, fault, forcibly stop, and monitor resources using scripts. Sometimes called an entry point.

base IP address
The first logical IP address, can be used as an administrative IP address.

entry point
See agent function.

floating IP address
See virtual IP address.

logical IP address
Any IP address assigned to a NIC.

NIC bonding
Combining two or more NICs to form a single logical NIC, which creates a fatter pipe.

operation
All agents have scripts that turn the resource on and off. Operations determine the action that the agent passes to the resource. See None operation, OnOff operation, and OnOnly operation.

None operation
For example the NIC resource. Also called persistent resource, this resource is always on. This kind of resource has no online and offline scripts, and only monitors a resource.

OnOff operation
For example the IP and Share agents--in fact most agents are OnOff. This resource has online and offline scripts. Often this type of resource does not appear in the types file because by default when a resource does not have this resource type defined, it is OnOff.

OnOnly operation
For example the NFS, FileOnOnly resources. This kind of resource has an online script, but not an offline one.

plumb
Term for enabling an IP address—used across all platforms in this guide.

test IP address

IP addresses to help determine the state of a link by sending out a ping probe to another NIC (on another system.) Requires a return ping to complete the test. Test IP addresses can be the same as base IP addresses.

virtual IP address

IP addresses that can move from one NIC to another or from one node to another. VCS fails over these IP address with your application. Sometimes called a floating IP address.

Index

B

C

D

E

F

H

I

K

L

M

N

O

P

R

S

T

V