

Veritas™ Cluster Server Bundled Agents Reference Guide

Linux

5.1 Service Pack 1

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: 5.1 SP1

Document version: 5.1.SP1.2

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

All support services will be delivered in accordance with your support agreement

and the then-current enterprise technical support policy.

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level

- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

Product guides are available on the software disc in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

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

docs@symantec.com

About Symantec Connect

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

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

Contents

Chapter 1	Introduction	
	Resources and their attributes	17
	Modifying agents and their resources	18
	Attributes	18
	Enabling debug log messages	19
Chapter 2	Storage agents	
	About the storage agents	21
	DiskGroup agent	22
	Dependencies	22
	Agent functions	23
	State definitions	24
	Attributes	25
	Resource type definition	29
	DiskGroup agent notes	29
	High availability fire drill	29
	Using volume sets	30
	Setting the noautoimport flag for a disk group	30
	Sample configurations	30
	DiskGroup resource configuration	30
	Debug log levels	31
	DiskGroupSnap agent	32
	Dependencies	32
	Agent functions	33
	State definitions	33
	Attributes	34
	DiskGroupSnap agent notes	34
	Configuring the SystemZones attribute for the fire drill service group	35
	Configuring the firedrill service group	35
	Adding the ReuseMntPt attribute to the ArgList attribute for the Mount agent type	35
	Configuration considerations	36
	Agent limitations	37
	Resource type definition	37

Sample configurations	38
Typical main.cf configuration	39
Oracle main.cf configuration	41
Debug log levels	44
DiskReservation agent	45
Agent functions	45
State definitions	45
Attributes	46
Resource type definition	48
DiskReservation agent notes	48
The DiskReservation agent does not reserve disks that have multiple paths	48
Configuring the MonitorTimeout attribute for more than three disks	48
Sample configurations	49
Configuration 1	49
Debug log levels	49
Volume agent	50
Dependencies	50
Agent functions	50
State definitions	50
Attributes	51
Resource type definition	52
Sample configuration	52
Debug log levels	52
VolumeSet agent	53
Dependencies	53
Agent functions	53
State definitions	54
Attributes	54
Resource type definition	54
Sample configurations	55
A configured VolumeSet that is dependent on a DiskGroup resource	55
Agent notes	55
Inaccessible volumes prevent the VolumeSet agent from coming online	55
Debug log levels	55
LVMLogicalVolume agent	56
Dependencies	56
Agent functions	56
State definitions	56
Attributes	57

Resource type definition	57
Sample configurations	57
Configuration 1	57
Configuration 2	58
Debug log levels	58
LVMVolumeGroup agent	59
Dependencies	59
Agent functions	59
State definitions	60
Attributes	60
Resource type definition	61
LVMVolumeGroup agent notes	61
Enabling volume group activation protection for Veritas Dynamic Multi-Pathing	61
Sample configurations	62
Linux configuration 1	62
Linux configuration 2	62
Linux configuration 3	63
Linux Configuration 4	63
Debug log levels	64
Mount agent	65
Dependencies	65
Agent functions	65
State definitions	68
Attributes	69
Resource type definition	75
Mount agent notes	76
High availability fire drill	76
VxFS file system lock	76
IMF usage notes	77
Enabling second level monitoring for the Mount agent	77
Sample configurations	78
Basic SCSI configuration	78
VxFS mount lock example	78
NFS mount example	78
Debug log levels	78

Chapter 3 Network agents

About the network agents	79
Agent comparisons	80
IP and NIC agents	80
IPMultiNIC and MultiNICA agents	80
802.1Q trunking	80

IP agent	81
High availability fire drill	81
Dependencies	81
Agent functions	82
State definitions	82
Attributes	83
Resource type definition	85
Sample configurations	86
Configuration 1	86
Configuration using specified NetMask	86
Debug log levels	86
NIC agent	87
High availability fire drill	87
Dependencies	87
Bonded network interfaces	88
Agent functions	88
State definitions	88
Attributes	89
Resource type definition	90
Monitoring bonded NICs	91
Setting Mii and miimon	91
Sample configurations	92
Configuration for using Mii	92
Configuration for using network hosts	92
IPv6 configuration	92
Debug log levels	93
IPMultiNIC agent	94
Dependencies	94
Agent functions	94
State definitions	95
Attributes	96
Resource type definition	98
Sample configuration: IPMultiNIC and MultiNICA	98
IPv4 configuration	98
IPv6 configuration	99
Mixed mode configuration—IPv4 and IPv6	100
Debug log levels	101
MultiNICA agent	102
Dependencies	102
IP Conservation Mode (ICM)	103
Configuration	103
Operation	103
Performance Mode (PM)	103

Configuration	103
Operation	104
Agent function	104
Attributes	104
Resource type definition	110
Sample configurations	110
MultiNICA and IPMultiNIC Performance Mode configuration ...	110
MultiNICA and IPMultiNIC IP Conservation Mode	
Configuration	111
IPv6 configuration	113
Mixed mode configuration—IPv4 and IPv6	113
Debug log levels	114
DNS agent	115
Dependencies	115
Agent functions	116
State definitions	117
Attributes	118
Resource type definition	122
DNS agent notes	122
High availability fire drill	123
Monitor scenarios	123
Sample Web server configuration	123
Secure DNS update for BIND 9	124
Setting up secure updates using TSIG keys for BIND 9	124
Sample configurations	125
Basic IPv6 configuration	125
IPv6 CNAME sample configuration	126
IPv4 A sample configuration	126
Debug log levels	126

Chapter 4 File share agents

About the file service agents	127
NFS agent	128
Dependencies	128
Agent functions	129
State definitions	129
Attributes	130
Resource type definition	131
NFS agent notes	131
Prerequisites for NFS lock recovery	131
Using NFSv4	131
Sample configurations	132
Debug log levels	132

NFSRestart agent	133
Dependencies	133
Agent functions	134
State definitions	135
Attributes	136
Resource type definition	137
NFSRestart agent notes	137
About high availability fire drill	137
Providing a fully qualified host name	137
Sample configurations	138
Basic agent configurations	138
Debug log levels	139
Share agent	140
Dependencies	140
Agent functions	141
State definitions	141
Attributes	142
Resource type definition	143
Share agent notes	143
High availability fire drill	143
Sample configurations	144
Debug log levels	144
About the Samba agents	145
The Samba agents	145
Before using the Samba agents	145
Supported versions	146
Notes for configuring the Samba agents	146
Enabling VCS to detect services started and stopped by smb	146
Configuring multiple SambaServer resources	146
Configuring Samba for non-standard configuration files or non- standard lock directories	147
SambaServer agent	148
Dependencies	148
Agent functions	148
State definitions	149
Attributes	150
Resource type definitions	152
Sample configurations	152
Debug log levels	152
SambaShare agent	153
Dependencies	153
Agent functions	153
State definitions	153

Attributes	154
Resource type definition	154
Sample configuration	155
Debug log levels	155
NetBios agent	156
Dependencies	156
Agent functions	157
State definitions	157
Attributes	158
Resource type definition	159
Sample configuration	160
Debug log levels	160

Chapter 5 Service and application agents

About the service and application agents	161
Apache Web server agent	162
Dependencies	162
Agent functions	163
State definitions	163
Attributes	164
Resource type definition	168
Apache Web server notes	169
Tasks to perform before you use the Apache Web server agent	169
About detecting application failure	170
About bringing an Apache Web server online outside of VCS control	170
About high Availability fire drill	170
Sample configurations	171
Running two versions of httpd for Linux	171
Basic IPv6 configuration	175
Application agent	176
High availability fire drill	176
Dependencies	176
Agent functions	178
State definitions	179
Attributes	180
Resource type definition	183
Application agent notes	184
Using Application agent with IMF	184
Sample configurations	184
Configuration 1	184
Configuration 2	185
Debug log levels	185

CoordPoint agent	186
Dependencies	186
Agent functions	186
State definitions	187
Attributes	187
Resource type definition	188
Notes for the CoordPoint agent	188
CoordPoint agent I/O fencing reporting activities	188
AutoStartList attribute	188
Sample configuration	189
Debug log levels	189
Process agent	190
High availability fire drill	190
Dependencies	190
Agent functions	191
State definitions	192
Attributes	192
Resource type definition	193
Sample configurations	194
Configuration for Linux	194
Debug log levels	194
ProcessOnOnly agent	195
Dependencies	195
Agent functions	195
State definitions	195
Attributes	196
Resource type definition	197
Sample configurations	197
Debug log levels	198

Chapter 6 Infrastructure and support agents

About the infrastructure and support agents	199
NotifierMngr agent	200
Dependency	200
Agent functions	200
State definitions	200
Attributes	201
Resource type definition	204
Sample configuration	205
Configuration	205
Debug log levels	206
Proxy agent	207
Dependencies	207

Agent functions	207
Attributes	208
Resource type definition	209
Sample configurations	209
Configuration 1	209
Configuration 2	209
Configuration 3	209
Debug log levels	210
Phantom agent	211
Dependencies	211
Agent functions	211
Resource type definition	211
Sample configurations	211
Configuration 1	211
Configuration 2	212
RemoteGroup agent	213
Dependency	213
Agent functions	214
State definitions	214
Attributes	215
Resource type definition	220
Debug log levels	220

Chapter 7 Testing agents

About the testing agents	221
ElifNone agent	222
Dependencies	222
Agent function	222
State definitions	222
Attributes	223
Resource type definition	223
Sample configuration	223
Debug log levels	223
FileNone agent	224
Dependencies	224
Agent functions	224
State definitions	224
Attribute	225
Resource type definition	225
Sample configuration	225
Debug log levels	225
FileOnOff agent	226
Dependencies	226

Agent functions	226
State definitions	226
Attribute	227
Resource type definition	227
Sample configuration	227
Debug log levels	227
FileOnOnly agent	228
Dependencies	228
Agent functions	228
State definitions	228
Attribute	229
Resource type definition	229
Sample configuration	229
Debug log levels	229
Glossary	231
Index	233

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 how agents work, refer to the *Veritas Cluster Server Administrator's Guide*.

Resources and their attributes

Resources are parts of a system. They are known by their types, for example: 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 Operations Manager, or the command line to dynamically modify the configuration of the resources managed by an agent.

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

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

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	<p>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 (_).</p> <p>A string can contain double quotes, but the quotes must be immediately preceded by a backslash. In a string, represent a backslash with two backslashes (\\).</p>
integer	<p>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: 2147483647.</p>

Table 1-1 Attribute data types

Data Type	Description
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{}.

Enabling debug log messages

To help troubleshoot agent issues, you can enable debug log messages in the agent framework as well as the agents.

To enable agent framework debug log messages:

```
hatype -modify agent_name LogDbg -add DBG_AGDEBUG DBG_AGINFO
DBG_AGTRACE
```

For example:

```
hatype -modify Mount LogDbg -add DBG_AGDEBUG DBG_AGINFO DBG_AGTRACE
```

To enable agent-specific debug log messages:

```
hatype -modify agent_name LogDbg -add debug_log_levels
```

For example:

```
hatype -modify Mount LogDbg -add DBG_1 DBG_2 DBG_3 DBG_4 DBG_5 DBG_6
```

Alternatively, you can also use the following command:

```
hatype -modify Mount LogDbg -add 1 2 3 4 5 6
```

Agent-specific debug log level information is specified in the agent's description. For example, for information about the Mount agent, see "[Debug log levels](#)" on page 86.

Storage agents

This chapter contains:

- [“About the storage agents”](#) on page 21
- [“DiskGroup agent”](#) on page 22
- [“DiskGroupSnap agent”](#) on page 32
- [“DiskReservation agent”](#) on page 45
- [“Volume agent”](#) on page 50
- [“VolumeSet agent”](#) on page 53
- [“LVMLogicalVolume agent”](#) on page 56
- [“LVMVolumeGroup agent”](#) on page 59
- [“Mount agent”](#) on page 65

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.

For important information on this agent, refer to:

“[DiskGroup agent notes](#)” on page 29

Dependencies

The DiskGroup resource does not depend on any other resources.

Figure 2-1 Sample service group that includes a DiskGroup resource



Agent functions

Online	Imports the disk group using the <code>vxdg</code> command.
Offline	Deports the disk group using the <code>vxdg</code> command.
Monitor	<p>Determines if the disk group is online or offline using the <code>vxdg</code> command. The Monitor function changes the value of the VxVM <code>noautoimport</code> flag from off to on. This action allows VCS to maintain control of importing the disk group. The monitor function uses following command to set the <code>noautoimport</code> flag to on.</p> <pre># vxdg -g disk_group set autoimport=no</pre>
Clean	Terminates all ongoing resource actions and takes the resource offline—forcibly when necessary.
Info	<p>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 <code>InfoInterval</code> timing to a value greater than 0.</p> <p>In the following example, the info agent function executes every 60 seconds:</p> <pre># haconf -makerw # hatype -modify DiskGroup InfoInterval 60</pre> <p>The command to retrieve information about the <code>DiskType</code> and <code>FreeSize</code> of the DiskGroup resource is:</p> <pre># hares -value diskgroupres ResourceInfo</pre> <p>Output includes:</p> <pre>DiskType sliced FreeSize 35354136</pre>

Action	Different action agent functions follow: <ul style="list-style-type: none">■ <code>license.vfd</code> Checks for valid Veritas Volume manager license—if one is not found use the <code>vxlicinst</code> utility to install a valid license key.■ <code>disk.vfd</code> 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.■ <code>udid.vfd</code> Checks the UDIDs (unique disk identifiers) 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.■ <code>verifyplex.vfd</code> 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.■ <code>volinuse</code> 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. One cause of this state is when I/O fencing is not configured—the cluster level attribute <code>UseFence</code> is not set to <code>"SCSI3"</code> but the <code>Reservation</code> attribute value is <code>"SCSI3"</code> .

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
MonitorReservation	If the value 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

Table 2-2 Optional attributes

Optional attributes	Description
PanicSystemOnDGLoss	<p>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.</p> <p>If the value of this attribute is 1, and the disk group becomes disabled, the node panics.</p> <p>If the value of this attribute is 1, and the Monitor agent function (entry point) hangs a consecutive number of times per the value of the FaultOnMonitorTimeouts attribute, then the node panics.</p> <p>Note: System administrators may want to set a high value for FaultOnMonitorTimeout to increase system tolerance.</p> <p>If the value of the attribute is 0, and the disk group becomes disabled, the following occurs:</p> <ul style="list-style-type: none"> ■ If the cluster has I/O fencing enabled, the DiskGroup resource is marked <code>FAULTED</code>. 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 <code>FAULTED</code> 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 <code>ONLINE</code>. The resource is reported <code>ONLINE</code> so that it does not fail over, which ensures data integrity. <p>Note: The PanicSystemOnDGLoss attribute does not depend on the MonitorReservation attribute.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>

Table 2-2 Optional attributes

Optional attributes	Description
StartVolumes	<p>If value of this attribute is 1, the DiskGroup online function starts all volumes belonging to that disk group after importing the group.</p> <p>Note: With VxVM version 5.1.100.0 onwards, if the Veritas Volume Manager default autostartvolumes at system level is set to on, all the volumes of the disk group will be started as a part of the import disk group.</p> <p>Type and dimension: boolean-scalar Default: 1</p>
StopVolumes	<p>If value is 1, the DiskGroup offline function stops all volumes belonging to that disk group before it deports the disk group.</p> <p>Type and dimension: boolean-scalar Default: 1</p>
UmountVolumes	<p>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:</p> <ul style="list-style-type: none">■ 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. <p>Type and dimension: integer-scalar Default: 0</p>

Table 2-2 Optional attributes

Optional attributes	Description
Reservation	<p>Determines if you want to enable SCSI-3 reservation. This attribute can have one of the following three values:</p> <ul style="list-style-type: none"> ■ ClusterDefault—The disk group is imported with SCSI-3 reservation if the value of the cluster-level UseFence attribute is SCSI3. If the value of the cluster-level UseFence attribute is NONE, the disk group is imported without reservation. ■ SCSI3—The disk group is imported with SCSI-3 reservation if the value of the cluster-level UseFence attribute is SCSI3. ■ NONE—The disk group is imported without SCSI-3 reservation. <p>Type and dimension: string-scalar Default: ClusterDefault Example: "SCSI3"</p>

Table 2-3 Internal attributes

Required attribute	Description
tempUseFence	Do not use. For internal use only.
NumThreads	<p>Number of threads used within the agent process for managing resources. This number does not include threads used for other internal purposes.</p> <p>Do not modify this attribute for this agent.</p> <p>Setting this attribute to a higher value may result in agent function timeouts due to serialization of underlying commands.</p> <p>Type and dimension: static integer-scalar Default: 1</p>
DiskGroupType	Do not use. For internal use only.

Resource type definition

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

DiskGroup agent notes

The DiskGroup agent has the following notes:

- [“High availability fire drill”](#) on page 29
- [“Using volume sets”](#) on page 30
- [“Setting the noautoimport flag for a disk group”](#) on page 30

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 Administrator’s Guide*.

Using volume sets

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

See the Mount agent description for more information.

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.

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
```

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
```

Sample configurations

DiskGroup resource configuration

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

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

Debug log levels

The DiskGroup agent uses the following debug log levels:

DBG_1, DBG_4, DBG_5

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. The agent also supports SCSI-3 fencing.

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

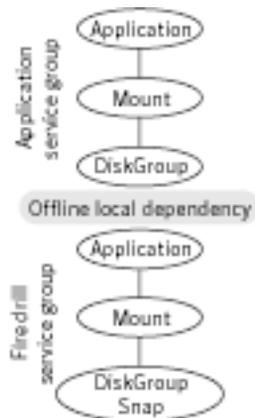
For important information about this agent, refer to:

“[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's resource has an offline local dependency on the application's service group. The offline local dependency is to make sure the fire drill service group and the application service group are not online at the same site at the same time.

Figure 2-2 Sample service group that includes 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.

State definitions

ONLINE	The DiskGroupSnap resource functions normally.
OFFLINE	The DiskGroupSnap resource is not running.
UNKNOWN	A configuration error exists.
FAULTED	The DiskGroupSnap resource is taken offline unexpectedly outside of VCS control.

Attributes

Table 2-4 Required attributes

Required attribute	Description
TargetResName	<p>The name of the DiskGroup resource from the application service group.</p> <p>Type-dimension: string-scalar</p> <p>Example: "dgres"</p>
FDSiteName	<p>At a site, this is the unique VxVM site name tag for the fire drill disks. You can run the fire drill in the following configurations:</p> <ul style="list-style-type: none"> ■ In the Gold configuration, a site has a dedicated set of fire drill disks. In Figure 2-4, the disaster recovery site uses a Gold configuration. ■ In the Bronze configuration, a site uses its data disks as fire drill disks. In Figure 2-4, the primary site uses a Bronze configuration. <p>Type and dimension: string-scalar</p> <p>Example:</p> <p>The value for the FDSiteName attribute for the configuration for Figure 2-4 is:</p> <pre>"FDSiteName@Node_A = pri" "FDSiteName@Node_B = pri" "FDSiteName@Node_C = dr_fd" "FDSiteName@Node_D = dr_fd"</pre>

DiskGroupSnap agent notes

The DiskGroupSnap agent has the following notes:

- [“Configuring the SystemZones attribute for the fire drill service group”](#) on page 35
- [“Configuring the firedrill service group”](#) on page 35
- [“Adding the ReuseMntPt attribute to the ArgList attribute for the Mount agent type”](#) on page 35
- [“Configuration considerations”](#) on page 36
- [“Agent limitations”](#) on page 37

Configuring the SystemZones attribute for the fire drill service group

You must assign the local system values to the SystemZones attribute of the application's service group. You set these values so that the service group fails over in the same zone before it tries to fail over across zones. For more information about campus cluster setup, refer to the *Veritas Cluster Server Administrator's Guide*.

For example, you set up the service group's SystemZones attribute for two zones: 0 and 1. You want the service group on Node_A and Node_B to fail over between the two nodes before it comes up on Node_C and Node_D. The application and its fire drill service group both have the following values for the SystemZones attribute:

```
SystemZones = { Node_A = 0, Node_B = 0, Node_C = 1, Node_D = 1 }
```

Configuring the firedrill service group

In the firedrill service group, the application-level resources (for example, process resources, application resources, or Oracle resources, and so on) can have the same attribute values in the firedrill service group and the application service group. The reuse of the same values for the attributes can result in VCS reporting the wrong resources as online.

Set the FireDrill type-level attribute to 1 for those types. For example, if the Oracle and Listener resources are configured identically, set the FireDrill attribute for Oracle and Listener to 1:

```
haconf -makerw
hatype -modify Oracle FireDrill 1
hatype -modify Listener FireDrill 1
haconf -dump -makero
```

Adding the ReuseMntPt attribute to the ArgList attribute for the Mount agent type

If you plan to use a Mount resource in a firedrill service group, you must add the ReuseMntPt attribute to ArgList and set its value to 1.

To add the ReuseMntPt attribute to the ArgList attribute and set its value to 1

- 1 Make the configuration read and write.
haconf -makerw
- 2 Add the ReuseMntPt attribute to the ArgList attribute.
hatype -modify Mount ArgList -add ReuseMntPt
- 3 Change the value of the ReuseMntPt attribute to 1 for the firedrill's Mount resource.

- ```
hares -modify firedrill_mount_resource_name ReuseMntPt 1
```
- 4 Change the value of the ReuseMntPt attribute to 1 for the original Mount resource.
- ```
# hares -modify original_mount_resource_name ReuseMntPt 1
```
- 5 Make the configuration read only.
- ```
haconf -dump -makero
```

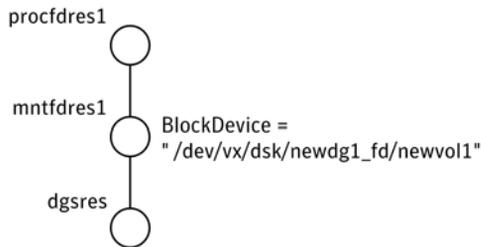
## Configuration considerations

Keep the following recommendations in mind:

- You must install Veritas Volume Manager 5.1 or later with the FMR license and the Site Awareness license.
- Do not bring the DiskGroupSnap resource online in the SystemZone where the application service group is online.
- Make sure that the firedrill service group and the application service group both use the same values for the SystemZones attribute.
- Do not use Volume resources in the firedrill service group. The DiskGroupSnap agent internally uses the `vxvol` command to start all the volumes in the firedrill disk group.
- In large setups, you may need to tweak the various timer values so that the timers do not time out while waiting for VxVM commands to complete. The timers you need to tweak are the OfflineTimeout for the DiskGroupSnap resource and MonitorInterval and ActionTimeout for the associated DiskGroup resource, for example:

```
haconf -makerw
hares -override dgsres OfflineTimeout
hares -modify dgsres OfflineTimeout 600
hares -override dgres MonitorInterval
hares -modify dgres MonitorInterval 1200 (this has to be twice
the value intended for ActionTimeout below)
hares -override dgres ActionTimeout
hares -modify dgres ActionTimeout 600
haconf -dump -makero
```

- When you create the firedrill service group, in general use the same attribute values that you use in the application service group. The BlockDevice attribute of the Mount resource changes between the application service group and the firedrill service group. In the BlockDevice path, you must append an `_fd` to the disk group name portion, for example, `/dev/vx/dsk/newdg1/newvol1` becomes `/dev/vx/dsk/newdg1_fd/newvol1`. [Figure 2-3](#) shows the changes to resource values for the firedrill 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 DiskGroupSnap agent does not support Volume Sets.
- The DiskGroupSnap agent cannot be used in a Storage Foundation RAC environment.
- 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.
- Take the firedrill service group offline before you shut down VCS on any node. If you fail to take the firedrill service group offline before you shut down VCS, you must manually reattach the fire drill site to the disk group to continue to perform fire drills.
- Use the enclosures that have the ASL/APM libraries that are supported in the Veritas Volume Manager. To view the supported enclosures, use the `vxddladm listsupport` command.

## 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

In [Figure 2-4](#), the Primary site is in the Bronze configuration and the Disaster recovery site is in a Gold configuration.

Since the Primary site does not have dedicated fire drill disks, it is in a Bronze configuration. In the Bronze configuration, you re-purpose the mirror disks in the disaster recovery site to serve as fire drill test disks. The drawback with the Bronze configuration is that if a disk failure occurs when the fire drill is online at the Primary site, it results in a site failure.

The `FDSiteName` value in a bronze configuration is the VxVM site name. For this configuration, the `FDSiteName` attribute values for the nodes at the Primary site follow:

```
FDSiteName@Node_A = pri
FDSiteName@Node_B = pri
```

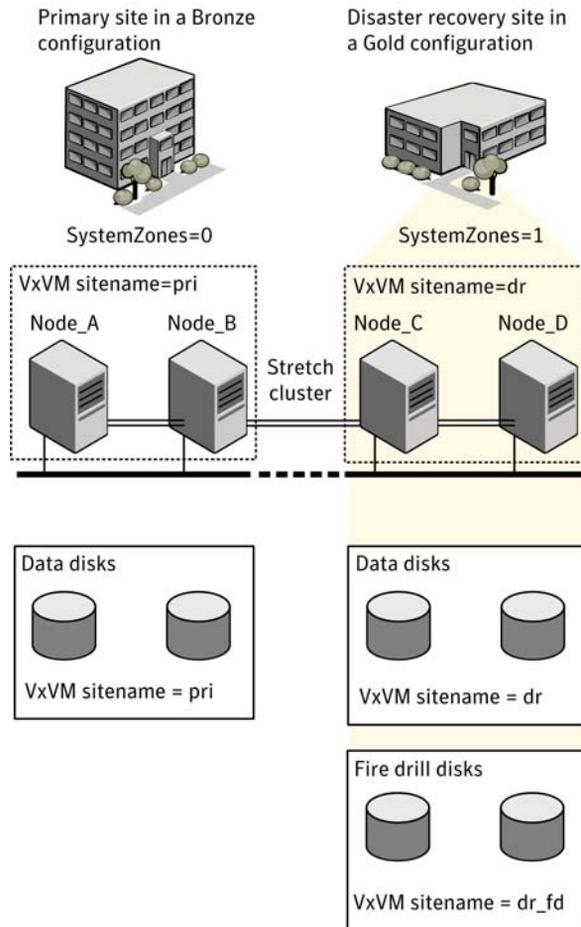
The Disaster Recovery site is in a Gold configuration as it has dedicated fire drill disks at the site. For the `FDSiteName` attribute, use the VxVM site tag given to the fire drill disks. For this configuration, the `FDSiteName` attribute values for the nodes at the Disaster recovery site follow:

```
FDSiteName@Node_C = dr_fd
FDSiteName@Node_D = dr_fd
```

Set values for the `SystemZones` attribute to zero for `Node_A` and `Node_B`, and one for `Node_C` and `Node_D`. For example:

```
SystemZones = { Node_A = 0, Node_B = 0, Node_C = 1, Node_D = 1 }
```

**Figure 2-4** Primary site with the Bronze configuration and the disaster recovery site with the Gold configuration



### Typical main.cf configuration

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

```
group dgfdsg (
 SystemList = { Node_A = 0, Node_B = 1, Node_C = 2, Node_D = 3 }
 SystemZones = { Node_A = 0, Node_B = 0, Node_C = 1, Node_D = 1 }
)

DiskGroupSnap dgsres (
 TargetResName = dgres
```

```

 FDSiteName @Node_A = pri
 FDSiteName @Node_B = pri
 FDSiteName @Node_C = dr_fd
 FDSiteName @Node_D = dr_fd
)

Mount mntfdres1 (
 MountPoint = "/dgsfs1"
 BlockDevice = "/dev/vx/dsk/newdg1_fd/newvol1"
 FSType = vxfs
 FsckOpt = "-y"
 ReuseMntPt = 1
)

Mount mntfdres2 (
 MountPoint = "/dgsfs2"
 BlockDevice = "/dev/vx/dsk/newdg1_fd/newvol2"
 FSType = vxfs
 FsckOpt = "-y"
 ReuseMntPt = 1
)

Process procfres1 (
 PathName = "/usr/bin/ksh"
 Arguments = "/scrib.sh /dgsfs1"
)

Process procfres2 (
 PathName = "/usr/bin/ksh"
 Arguments = "/scrib.sh /dgsfs2"
)

requires group dgsg offline local
mntfdres1 requires dgsres
mntfdres2 requires dgsres
procfres1 requires mntfdres1
procfres2 requires mntfdres2

```

The application's service group (the actual service group) follows:

```

group dgsg (
 SystemList = { Node_A = 0, Node_B = 1, Node_C = 2, Node_D = 3 }
 SystemZones = { Node_A = 0, Node_B = 0, Node_C = 1, Node_D = 1 }
)

DiskGroup dgres (
 DiskGroup = newdg1
)

Mount mntres1 (
 MountPoint = "/dgsfs1"
 BlockDevice = "/dev/vx/dsk/newdg1/newvol1"
 FSType = vxfs

```

```
FsckOpt = "-y"
ReuseMntPt = 1
)

Mount mntres2 (
 MountPoint = "/dgsfs2"
 BlockDevice = "/dev/vx/dsk/newdg1/newvol2"
 FSType = vxfs
 FsckOpt = "-y"
 ReuseMntPt = 1
)

Process proces1 (
 PathName = "/usr/bin/ksh"
 Arguments = "/scrib.sh /dgsfs1"
)

Process proces2 (
 PathName = "/usr/bin/ksh"
 Arguments = "/scrib.sh /dgsfs2"
)

mntres1 requires dgres
mntres2 requires dgres
proces1 requires mntres1
proces2 requires mntres2
```

## Oracle main.cf configuration

The following Oracle configuration has been simplified for presentation within this guide. Note that *NIC0* represents the NIC's name.

```
group fd_oragrp (
 SystemList = { Node_A = 0, Node_B = 1 }
 AutoStart = 0
 SystemZones = { Node_A = 0, Node_B = 1 }
)

DiskGroupSnap dgres (
 FDSiteName @Node_A = siteA
 FDSiteName @Node_B = siteB
 TargetResName = oradg_res
)

IP fd_oraip (
 Device = NIC0
 Address = "10.198.95.191"
)

Mount fd_archmnt (
 FsckOpt = "-y"
```

```
 ReuseMntPt = 1
 BlockDevice = "/dev/vx/dsk/oradg_fd/archive_vol"
 MountPoint = "/ora_archive"
 FSType = vxfs
)

Mount fd_datamnt (
 FsockOpt = "-y"
 ReuseMntPt = 1
 BlockDevice = "/dev/vx/dsk/oradg_fd/data_vol"
 MountPoint = "/ora_data"
 FSType = vxfs
)

NIC fd_oranic (
 Device = NIC0
)

Netlsnr fd_LSNR (
 Home = "/opt/oracle/ora_home"
 Owner = oracle
)

Oracle fd_Ora_01 (
 Owner = oracle
 Home = "/opt/oracle/ora_home"
 Sid = Ora_01
)

requires group oragrp offline local
fd_LSNR requires fd_Ora_01
fd_LSNR requires fd_oraip
fd_Ora_01 requires fd_archmnt
fd_Ora_01 requires fd_datamnt
fd_archmnt requires dgres
fd_datamnt requires dgres
fd_oraip requires fd_oranic

group oragrp (
 SystemList = { Node_A = 0, Node_B = 1 }
 AutoStartList = { Node_A, Node_B }
 SystemZones = { Node_A = 0, Node_B = 1 }
)

DiskGroup oradg_res (
 DiskGroup = oradg
)

IP Node_A4vip (
 Device = NIC0
 Address = "10.198.95.192"
)
```

```
Mount arch_mnt (
 FsckOpt = "-y"
 ReuseMntPt = 1
 BlockDevice = "/dev/vx/dsk/oradg/archive_vol"
 MountPoint = "/ora_archive"
 FSType = vxfs
)
```

```
Mount data_mnt (
 FsckOpt = "-y"
 ReuseMntPt = 1
 BlockDevice = "/dev/vx/dsk/oradg/data_vol"
 MountPoint = "/ora_data"
 FSType = vxfs
)
```

```
NIC nic_Node_A4vip (
 Device = NIC0
)
```

```
Netlsnr LSNR (
 Home = "/opt/oracle/ora_home"
 Owner = oracle
)
```

```
Oracle Ora_01 (
 Owner = oracle
 Home = "/opt/oracle/ora_home"
 Sid = Ora_01
)
```

```
Volume arch_vol (
 Volume = archive_vol
 DiskGroup = oradg
)
```

```
Volume data_vol (
 Volume = data_vol
 DiskGroup = oradg
)
```

```
LSNR requires Ora_01
LSNR requires Node_A4vip
Ora_01 requires arch_mnt
Ora_01 requires data_mnt
arch_mnt requires arch_vol
arch_vol requires oradg_res
data_mnt requires data_vol
data_vol requires oradg_res
Node_A4vip requires nic_Node_A4vip
```

## Debug log levels

The DiskGroupSnap agent uses the following debug log levels:

DBG\_1

# DiskReservation agent

Reserves and monitors SCSI disks for a system, enabling a resource to go online on that system. This agent enables you to specify a list of raw disk devices, and reserve all or a percentage of accessible disks. The reservations prevent disk data corruption by restricting other nodes from accessing and writing to the disks. The DiskReservation agent supports all SCSI-II compliant disks.

An automatic probing feature allows systems to maintain reservations even when the disks or bus are reset. The optional FailFast feature minimizes data corruption in the event of a reservation conflict by causing the system to panic.

---

**Note:** The DiskReservation agent cannot be used to reserve disks that have multiple paths.

---

For important information on this agent, refer to:  
[“DiskReservation agent notes”](#) on page 48

## Agent functions

|         |                                                                                                 |
|---------|-------------------------------------------------------------------------------------------------|
| Online  | Brings the resource online after reserving all or a specified percentage of accessible disks.   |
| Offline | Releases reservations on reserved disks.                                                        |
| Monitor | Monitors the accessibility and reservation status of the reserved disks.                        |
| Clean   | Terminates all ongoing resource actions and takes the resource offline—forcibly when necessary. |

## State definitions

|         |                                                                                                                              |
|---------|------------------------------------------------------------------------------------------------------------------------------|
| ONLINE  | Indicates that the number of reserved disks is greater than or equal to the percentage specified in the resource definition. |
| OFFLINE | Disks are not reserved.                                                                                                      |
| UNKNOWN | Indicates that a problem exists with the configuration.                                                                      |

## Attributes

**Table 2-5** Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disks              | <p>A list of raw disk devices. Use the absolute or relative device path.</p> <p>The absolute or relative device path allows a maximum of 64 characters. The relative path is assumed to start from the <code>/dev</code> directory.</p> <p>The order of the disks in the list must be the same across all systems in the cluster, even if the same device has a different name on different systems.</p> <p><b>Note:</b> You must change this attribute before bringing a resource online. An online device must be taken offline before altering this attribute because disk reservation occurs during the process of bringing a resource online.</p> <p>Type and dimension: string-vector</p> <p>Example: "sdd"</p> |

**Table 2-6** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FailFast           | <p>If enabled, FailFast causes the system to panic when a reservation conflict is detected, reducing the chance of further data corruption.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> |

Table 2-6 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Percentage         | <p>Minimum percentage of configured disks that can be reserved before a resource can go online. The percentage must be greater than or equal to 51, and less than or equal to 100.</p> <p>If the value specified is less than 51, the percentage is set to 51.</p> <p>If the value specified is greater than 100, the percentage is set to 100.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 100</p>                                                                                                                                                                                               |
| ProbeInterval      | <p>Alters the periodicity (in seconds) of the automatic probe function that checks the reservation status of the disks. The value must be greater than or equal to three, and less than or equal to 15.</p> <p>If the value specified is less than 3, the interval is set to 3.</p> <p>If the value specified is greater than 15, the interval is set to 15.</p> <p>A lower value for ProbeInterval specifies more frequent probes and provides for quicker discovery of reservation conflicts. Symantec recommends a value is between 3 and 8.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 3</p> |

## Resource type definition

```
type DiskReservation (
 static str ArgList[] = { Disks, FailFast, Percentage,
 ProbeInterval }
 str Disks[]
 boolean FailFast = 0
 int Percentage = 100
 int ProbeInterval = 3
)
```

## DiskReservation agent notes

The DiskReservation agent has the following notes:

- [“The DiskReservation agent does not reserve disks that have multiple paths”](#) on page 48
- [“Configuring the MonitorTimeout attribute for more than three disks”](#) on page 48

### The DiskReservation agent does not reserve disks that have multiple paths

You cannot use the DiskReservation agent to reserve disks that have multiple paths. In case of Veritas Dynamic Multi-Pathing, the LVMVolumeGroup and the LVMLogicalVolume agents can be used without the DiskReservation agent. To ensure data protection on such a configuration, Symantec recommends the use of Veritas Volume Manager (VxVM) disk groups. Note that VxVM requires the use of SCSI-3 compliant disks.

### Configuring the MonitorTimeout attribute for more than three disks

The MonitorTimeout attribute’s setting of 60 is adequate for up to three disks. When configuring the MonitorTimeout attribute for more than three disks, use the following formula:

Set MonitorTimeout to be equal or greater than 15 times the total number of disks. (MonitorTimeout  $\geq$  15 \* Number of disks).

For example, if you have eight disks, MonitorTimeout is 120 or greater.

## Sample configurations

### Configuration 1

In this example, the DiskReservation agent reserves a disk. The disk is mounted with the Veritas File System.

```
system sysA

system sysB

group groupx (
 SystemList = { sysA, sysB }
 AutoStartList = { sysA }
)

DiskReservation diskres1 (
 Disks = { "/dev/sdc" }
 FailFast = 1
)

Mount mount (
 MountPoint = "/mnt/tmp"
 BlockDevice = "/dev/sdc1"
 FSType = vxfs
 MountOpt = rw
)

mount requires diskres1

// resource dependency tree
//
// group groupx
// {
// Mount mount
// {
// DiskReservation diskres1
// }
// }
// }
```

### Debug log levels

The DiskReservation agent uses the following debug log levels:

DBG\_1, DBG\_3, DBG\_4, DBG\_5

## Volume agent

The Volume agent brings online, takes offline, and monitors a Veritas Volume Manager (VxVM) volume. Use the agent to make a volume highly available.

---

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

---

## Dependencies

Volume resources depend on DiskGroup resources.

**Figure 2-5** Sample service group that includes a Volume resource



## Agent functions

|         |                                                                                                                                 |
|---------|---------------------------------------------------------------------------------------------------------------------------------|
| Online  | Uses the <code>vxrecover</code> command to start the volume.                                                                    |
| Offline | Uses the <code>vxvol</code> command to stop the volume.                                                                         |
| Monitor | Attempts to read a block from the raw device interface to the volume to determine if the volume is online, offline, or unknown. |
| 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 stopped unexpectedly and that I/O is not permitted.                                                           |
| UNKNOWN | Indicates that the agent could not determine the state of the resource or that the resource attributes are configured incorrectly. |

## Attributes

**Table 2-7** Required attributes

| Required attribute | Description                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| DiskGroup          | Name of the disk group that contains the volume.<br>Type and dimension: string-scalar<br>Example: "DG1 "                        |
| Volume             | Name of the volume from disk group specified in DiskGroup attribute.<br>Type and dimension: string-scalar<br>Example: "DG1Vol1" |

**Table 2-8** Internal attribute

| Internal 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.<br>Do not modify this attribute for this agent.<br>Setting this attribute to a higher value may result in agent function timeouts due to serialization of underlying commands.<br>Default: 1 |

## Resource type definition

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

## Sample configuration

```
Volume sharedg_vol3 (
 Volume = vol3
 DiskGroup = sharedg
)
```

## Debug log levels

The Volume agent uses the following debug log levels:  
DBG\_1, DBG\_3, DBG\_5

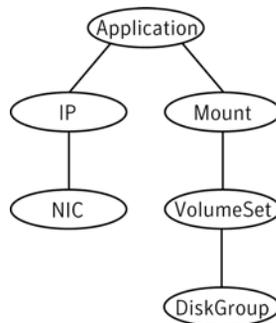
# VolumeSet agent

The VolumeSet agent brings online, takes offline, and monitors a Veritas Volume Manager (VxVM) volume set. Use the agent to make a volume set highly available.

## Dependencies

VolumeSet resources depend on DiskGroup resources.

**Figure 2-6** Sample service group that includes a VolumeSet resource



## Agent functions

|         |                                                                                                                                                            |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online  | Uses the <code>vxrecover</code> command to start the volume set.                                                                                           |
| Offline | Uses the <code>vxvol</code> command to stop the volume set.                                                                                                |
| Monitor | Attempts to read a block from the raw device interface to the volumes inside the volume set to determine if the volume set is online, offline, or unknown. |
| Clean   | Terminates all ongoing resource actions and takes the resource offline—forcibly when necessary.                                                            |

## State definitions

|         |                                                                                                                                    |
|---------|------------------------------------------------------------------------------------------------------------------------------------|
| ONLINE  | Indicates that all the volumes in the volume set are started and that I/O is permitted for all the volumes.                        |
| OFFLINE | Indicates that at least one of the volume is not started in the volume set and that I/O is not permitted for that volume.          |
| FAULTED | Indicates the volumes that are inside the volume set have stopped unexpectedly and that I/O is not permitted.                      |
| UNKNOWN | Indicates that the agent could not determine the state of the resource or that the resource attributes are configured incorrectly. |

## Attributes

**Table 2-9** Required attributes

| Required attribute | Description                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DiskGroup          | The name of the disk group that contains the volume set.<br>Type and dimension: string-scalar<br>Example: "DG1"                                             |
| VolumeSet          | The name of the volume set from the disk group that you specified in the DiskGroup attribute.<br>Type and dimension: string-scalar<br>Example: "DG1VolSet1" |

## Resource type definition

```
type VolumeSet (
 static str ArgList[] = { DiskGroup, VolumeSet }
 str VolumeSet
 str DiskGroup
)
```

## Sample configurations

This sections contains sample configurations for this agent.

### A configured VolumeSet that is dependent on a DiskGroup resource

The VolumeSet's `shared_vset3` resource is configured and is dependent on DiskGroup resource with a shared diskgroup.

```
VolumeSet sharedg_vset3 (
 VolumeSet = vset3
 DiskGroup = sharedg
)
```

## Agent notes

This sections contains notes about this agent.

### Inaccessible volumes prevent the VolumeSet agent from coming online

The VolumeSet agent does not come online if any volume is inaccessible in its volume set.

#### To remove a volume from volume set

- ◆ Enter the following commands to remove a volume from a volume set mounted on *mountpoint*.

```
fsvoladm remove mountpoint volume_name
vxvset -g diskgroup rmvol volumeset volume_name
```

## Debug log levels

The VolumeSet agent uses the following debug log levels:

DBG\_1, DBG\_4

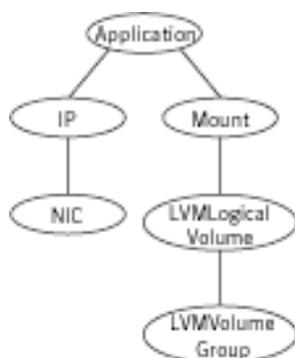
## LVMLogicalVolume agent

The LVMLogicalVolume agent brings online, takes offline, and monitors a Logical Volume Manager (LVM2) volume. This agent uses LVM2 commands. You can use this agent to make volume groups and logical volumes highly available and to monitor them.

### Dependencies

LVMLogicalVolume resources depend on LVMVolumeGroup resources.

**Figure 2-7** Sample service group that includes a LVMLogicalVolume resource



### Agent functions

|         |                                                                                                               |
|---------|---------------------------------------------------------------------------------------------------------------|
| Online  | Starts the volume using the <code>lvchange</code> command.                                                    |
| Offline | Stops the volume using the <code>lvchange</code> 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 I/O is not permitted.                                        |
| UNKNOWN | Indicates that the agent could not determine the state of the resource or that the resource attributes are invalid. |

## Attributes

**Table 2-10** Required attributes

| Required attribute | Description                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LogicalVolume      | Name of the volume that is configured with Logical Volume Manager (LVM2).<br>Type and dimension: string-scalar<br>Example: "volume1"                                       |
| VolumeGroup        | Name of the volume group that is configured with Logical Volume Manager (LVM2), which contains the volume.<br>Type and dimension: string-scalar<br>Example: "volumegroup1" |

## Resource type definition

```
type LVMLogicalVolume (
 static str ArgList[] = { LogicalVolume, VolumeGroup }
 str LogicalVolume
 str VolumeGroup
)
```

## Sample configurations

### Configuration 1

In this example, /dev/sdc and /dev/sdd are the disks where the volume group testvg\_1 is created.

```
LVMLogicalVolume lv011 (
 LogicalVolume = testvol_1
 VolumeGroup = testvg_1
)
```

```
LVMVolumeGroup lvg1 (
 VolumeGroup = testvg_1
)

DiskReservation dr1 (
 Disks = { "/dev/sdc", "/dev/sdd" }
)

lv011 requires lvg1
lvg1 requires dr1
```

## Configuration 2

In this example, you use the DiskReservation resource to verify that disks are available only to one system. The LVM2 logical volumes on the LVM2 volume groups that are imported are started if the reservation is confirmed. The logical volumes can then be mounted at a mount point.

```
DiskReservation dr_cde (
 Disks = { "/dev/sdc", "/dev/sdd", "/dev/sde" }
)

Mount mnt_lvmvol01 (
 MountPoint = "/mnt/lvmvol01"
 BlockDevice = "/dev/mapper/lvmvg01-lvmvol01"
 FSType = "reiserfs"
 FsckOpt = "-y"
)

LVMLogicalVolume lvmvol01 (
 LogicalVolume = lvmvol01
 VolumeGroup = lvmvg01
)

LVMVolumeGroup lvmvg01 (
 VolumeGroup = lvmvg01
)

mnt_lvmvol01 requires lvmvol01
lvmvol01 requires lvmvg01
lvmvg01 requires dr_cde
```

## Debug log levels

The LVMLogicalVolume agent uses the following debug log levels:

DBG\_1, DBG\_3, DBG\_5

# LVMVolumeGroup agent

The LVMVolumeGroup agent brings online, takes offline, and monitors a Logical Volume Manager (LVM2) volume group. This agent uses LVM2 commands. You can use this agent to make volume groups and logical volumes highly available and to monitor them.

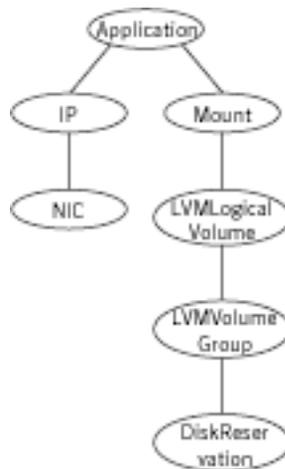
This agent supports Veritas Dynamic Multi-Pathing.

## Dependencies

No fixed dependencies exist for LVMVolumeGroup Agent. When you create a volume group on disks with single path, Symantec recommends that you use the DiskReservation agent.

With Veritas Dynamic Multi-Pathing, do not use the DiskReservation agent.

**Figure 2-8** Sample service group for a LVMVolumeGroup resource



## Agent functions

|         |                                                                                               |
|---------|-----------------------------------------------------------------------------------------------|
| Online  | Imports the volume group using the <code>vgimport</code> command.                             |
| Offline | Exports the volume group using the <code>vgexport</code> command.                             |
| Monitor | Determines if the volume group is online or offline using the <code>vgdisplay</code> command. |

|        |                                                                                                                                                                                                                                                       |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Clean  | Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.                                                                                                                                                      |
| Action | <ul style="list-style-type: none"> <li>■ volinuse<br/>Checks if the specified volume is mounted outside of VCS control or not. In case it is mounted outside of VCS control, the action returns "FAILURE", otherwise it returns "SUCCESS".</li> </ul> |

## State definitions

|         |                                                                                                                       |
|---------|-----------------------------------------------------------------------------------------------------------------------|
| ONLINE  | Indicates that the volume group is imported.                                                                          |
| OFFLINE | Indicates that the volume group is not imported.                                                                      |
| UNKNOWN | Indicates that a problem exists either with the configuration or the ability to determine the status of the resource. |

## Attributes

**Table 2-11** Required attributes

| Required attribute | Description                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VolumeGroup        | <p>The name of the volume group that is configured with Logical Volume Manager (LVM2) that contains the volume.</p> <p>Type and dimension: string-scalar</p> <p>Example: "volumegroup1"</p> |

**Table 2-12** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| StartVolumes       | <p>If the value of this attribute is 1, the LVMVolumeGroup online function imports the group. It then starts all the volumes that belong to that volume group.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> |

**Table 2-12** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EnableLVMTagging   | <p>If the value of this attribute is 1, the LVMVolumeGroup online function associates an LVM tag with the volume group and then imports the group. While taking the resource offline, it removes the associated tags.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> |

## Resource type definition

```

type LVMVolumeGroup (
 static keylist SupportedActions = { volinuse }
 static str ArgList[] = { VolumeGroup, StartVolumes,
 EnableLVMTagging }
 str VolumeGroup
 boolean StartVolumes = 0
 boolean EnableLVMTagging = 0
)

```

## LVMVolumeGroup agent notes

The LVMVolumeGroup agent has the following notes:

- [“Enabling volume group activation protection for Veritas Dynamic Multi-Pathing”](#) on page 61

### Enabling volume group activation protection for Veritas Dynamic Multi-Pathing

On each node in the cluster, perform the following procedure to enable activation protection for volume groups on Red Hat and SUSE systems.

#### To enable volume group activation protection

- 1 On each node in the cluster, edit `/etc/lvm/lvm.conf`, and add the following line:

```
tags { hosttags = 1 }
```
- 2 On each node in the cluster, create the file `lvm_`uname -n`.conf` in the `/etc/lvm/` directory.
- 3 Add the following line to the file you created in step 2:

```
activation { volume_list="@node" }
```

where `node` is the value of the `uname -n` command.

### Notes for volume group activation protection

Review the following notes if you plan to use volume group activation protection:

- Setting `hosttag=0` in `/etc/lvm/lvm.conf` disables the lvm tagging for all the volume groups.
- If the system's LVM configuration (`/etc/lvm/lvm.conf`) is modified to support LVM tagging (`hosttags = 1` in the `lvm.conf` file), make sure that the value of `EnableLVMTagging` is 1 for all the configured `LVMVolumeGroup` resources.

Another way to disable the LVM tagging for a particular volume group is to add that volume group name in a `volume_list` activation configuration (`/etc/lvm/lvm_`uname -n`.conf`).

For example, if you do not want LVM tagging to be enforced for the `vg11` volume group, then add the following line in `/etc/lvm/lvm_`uname -n`.conf`.

```
activation { volume_list=["vg11", "@node"] }
```

where "node" is the value of "uname -n" command.

Using this configuration instructs the agent to not use tagging while activating the volume group `vg11`. The rest of the volume groups require tagging however.

## Sample configurations

### Linux configuration 1

In this example, `/dev/sdc` and `/dev/sdd` are the disks with a single path where the volume group `testvg_1` is created.

```
LVMVolumeGroup lvg1 (
 VolumeGroup = testvg_1
)

DiskReservation dr1 (
 Disks = { "/dev/sdc", "/dev/sdd" }
)
```

```
lvg1 requires dr1
```

### Linux configuration 2

In this example, the `DiskReservation` resource is used to verify that disks are available only to one system. All LVM2 logical volumes on the LVM2 volume groups that are imported are started if the reservation is confirmed. You can then mount the logical volumes at a mount point.

```
DiskReservation dr_cde (
 Disks = { "/dev/sdc", "/dev/sdd", "/dev/sde" }
)

Mount mnt_lvmvol01 (
 MountPoint = "/mnt/lvmvol01"
 BlockDevice = "/dev/mapper/lvmvg01-lvmvol01"
 FSType = "reiserfs"
 FsckOpt = "-y"
)

LVMVolumeGroup lvmvg01 (
 VolumeGroup = lvmvg01
 StartVolumes = 1
)

mnt_lvmvol01 requires lvmvg01
lvmvg01 requires dr_cde
```

### Linux configuration 3

In this example, the volume group testvg\_1 is created on disks that have multiple paths and have LVM tagging enabled.

```
LVMVolumeGroup lvg1 (
 VolumeGroup = testvg_1
 EnableLVMTagging = 1
)
```

### Linux Configuration 4

In this example, the volume groups testvg\_1 and testvg\_2 are created on disks that have multiple paths. LVM tagging is enabled for testvg\_1 and disabled for test\_vg2.

```
LVMVolumeGroup lvg1 (
 VolumeGroup = testvg_1
 EnableLVMTagging = 1
)

LVMVolumeGroup lvg2 (
 VolumeGroup = testvg_2
)
```

Here the `/etc/lvm/lvm_`uname -n`.conf` file contains:

```
activation { volume_list=["testvg_2", "@node"] }
```

where node is the value of `uname -n` command.

## Debug log levels

The LVMVolumeGroup agent uses the following debug log levels:

DBG\_1, DBG\_3, DBG\_5

## 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 client mount points highly available. This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

---

**Note:** Intelligent Monitoring Framework for mounts is supported only for the following mount types: VxFS, CFS, and NFS.

---

For important information about this agent, refer to:  
“[Mount agent notes](#)” on page 76

## Dependencies

The Mount resource does not depend on any other resources.

**Figure 2-9** Sample service group that includes 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 before attempting to mount the file system again.

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.<br><br>If IMF is enabled for the Mount agent, the resource is monitored asynchronously and any change in the resource state is immediately sent to VCS for appropriate action.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| imf_init            | Initializes the agent to interface with the asynchronous monitoring framework (AMF) kernel driver. This function runs when the agent starts up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| imf_getnotification | Waits for notification about resource state changes. This function runs after the agent initializes with the AMF kernel driver. The agent continuously waits for notification and takes action on the resource upon notification.                                                                                                                                                                                                                                                                                                                                                                                                           |
| imf_register        | Registers the resource entities, which the agent must monitor, with the AMF kernel driver. This function runs for each resource after the resource goes into steady state (online or offline).                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Clean               | Unmounts the mounted file system forcefully.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Info                | The Mount info agent function executes the command:<br><pre>df -h <i>mount_point</i></pre> The output displays Mount resource information:<br><pre>Size Used Avail Use%</pre> 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:<br><pre>haconf -makerw<br/>hatype -modify Mount InfoInterval 60</pre> The command to retrieve information about the Mount resource is:<br><pre>hares -value mountres ResourceInfo</pre> <b>Output includes:</b><br><pre>Size 2097152<br/>Used 139484<br/>Available 1835332<br/>Used% 8%</pre> |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Offline             | Unmounts the mounted file system gracefully.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Monitor             | <p>Determines if the file system is mounted.</p> <p>If IMF is enabled for the Mount agent, the resource is monitored asynchronously and any change in the resource state is immediately sent to VCS for appropriate action.</p>                                                                                                                                                                                                                                                                                                                                                                                                                          |
| imf_init            | Initializes the agent to interface with the asynchronous monitoring framework (AMF) kernel driver. This function runs when the agent starts up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| imf_getnotification | Waits for notification about resource state changes. This function runs after the agent initializes with the AMF kernel driver. The agent continuously waits for notification and takes action on the resource upon notification.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| imf_register        | Registers the resource entities, which the agent must monitor, with the AMF kernel driver. This function runs for each resource after the resource goes into steady state (online or offline).                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Clean               | Unmounts the mounted file system forcefully.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Info                | <p>The Mount info agent function executes the command:</p> <pre>df -h <i>mount_point</i></pre> <p>The output displays Mount resource information:</p> <pre>Size Used Avail Use%</pre> <p>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:</p> <pre>haconf -makerw<br/>hatype -modify Mount InfoInterval 60</pre> <p>The command to retrieve information about the Mount resource is:</p> <pre>hares -value mountres ResourceInfo</pre> <p>Output includes:</p> <pre>Size 2097152<br/>Used 139484<br/>Available 1835332<br/>Used% 8%</pre> |

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action                    | <ul style="list-style-type: none"><li>■ <code>chgmntlock</code><br/>Resets the VxFS file system lock to a VCS-defined lock.</li><li>■ <code>mountpoint.vfd</code><br/>Checks if the specified mount point exists on the offline node. If it fails and you request that VCS fixes it, it creates the mount point directory using <code>mkdir</code> command.</li><li>■ <code>mounted.vfd</code><br/>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.</li><li>■ <code>vxfslic.vfd</code><br/>Checks for valid Veritas File System (VxFS) licenses. If it fails, you need to update the license for VxFS.</li><li>■ <code>mountentry.vfd</code><br/>Checks that the mount point is not listed in auto file system tables (for example, <code>/etc/fstab</code>).<br/>If this action fails, you need to remove the mount point from auto file system tables.</li></ul> |
| <code>attr_changed</code> | Unlocks the mounts when you change the value of the <code>VxFSMountLock</code> attribute from 1 to 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## State definitions

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

## Attributes

**Table 2-13** Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BlockDevice        | <p>Block device for mount point.</p> <p>For LVM2, use the actual mapper path to the volume.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ <code>"/dev/vx/dsk/myvcs_dg/myvol"</code></li> <li>■ LVM2 example</li> </ul> <pre># ls -la /dev/ora_vg/ora_vol lrwxrwxrwx 1 root root 26 Apr 17 04:48 /dev/ora_vg/ora_vol -&gt; /dev/mapper/ora_vg- ora_vol</pre> <p>Use the path <code>/dev/mapper/ora_vg-ora_vol</code> for the BlockDevice attribute.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| FckOpt             | <p>Mandatory for non-NFS mounts.</p> <p>Use this attribute to specify options for the <code>fsck</code> command. You must correctly set this attribute for local mounts. If the mount process fails, the <code>fsck</code> command is executed with the specified options before it attempts to remount the block device. Its value must include either <code>-y</code> or <code>-n</code>. Refer to the <code>fsck</code> manual page for more information.</p> <p>The <code>-y</code> argument enables the VxFS file systems to perform a log replay before a full <code>fsck</code> operation.</p> <p>For NFS mounts, the value of this attribute is not applicable and is ignored.</p> <p>Type and dimension: string-scalar</p> <p>VxFS example: <code>-y</code></p> <p><b>Note:</b> When you use the command line, add the <code>%</code> sign to escape <code>'</code>.<br/>For example: <code>hares -modify MntRes FckOpt %-y</code></p> |
| FSType             | <p>Type of file system.</p> <p>Supports <code>vxfs</code>, <code>bind</code>, <code>ext2</code>, <code>ext3</code>, <code>nfs</code>, or <code>reiserfs</code>.</p> <p>Type and dimension: string-scalar</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 2-13** Required attributes

| Required attribute | Description                                                                                |
|--------------------|--------------------------------------------------------------------------------------------|
| MountPoint         | Directory for mount point.<br>Type and dimension: string-scalar<br>Example: "/mnt/apache1" |

**Table 2-14** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CkptUmount         | If the value of this attribute is 1, this attribute automatically unmounts VxFS Storage Checkpoints when the file system is unmounted.<br>If the value of this attribute is 0, and Storage Checkpoints are mounted, then failover does not occur.<br>Type and dimension: boolean-scalar<br>Default: 1 |
| MountOpt           | Options for the <code>mount</code> command. Refer to the <code>mount</code> manual page for more information.<br>Do not specify <code>-o</code> in the MountOpt field.<br>Type and dimension: string-scalar<br>Example: "rw"                                                                          |

Table 2-14 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VxFSMountLock      | <p>This attribute is applicable to Veritas (VxFS) file systems. It controls the agent's use of the locking feature provided by vxfs to prevent accident unmounts.</p> <p>If the value of this attribute is 0, the agent does not lock the mount point when the resource is brought online. It does not monitor the status of the lock when the resource is online. No warnings appear if the mount has been locked with a key different than "VCS".</p> <p>If the value of this attribute is 1, during online, the agent uses the key "VCS" to lock the mount point. The monitor agent function monitors the locks during every cycle.</p> <ul style="list-style-type: none"><li>■ If the mountpoint is not locked, the agent locks it.</li><li>■ If the mountpoint is already locked with a key other than "VCS", the agent logs a warning. It then requests that you run the Chgmntlock action agent function.</li></ul> <p>During offline, the agent, as required, unlocks using whatever key needed.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> |
| SecondLevelMonitor | <p>This attribute has been deprecated.</p> <p>Instead of this attribute, use the LevelTwoMonitorFreq attribute. For more information, see <a href="#">“Enabling second level monitoring for the Mount agent”</a> on page 77.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| SecondLevelTimeout | <p>This attribute has been deprecated.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Table 2-14** Optional attributes

| Optional attribute  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SnapUnmount         | <p>If the value of this attribute is 1, this attribute automatically unmounts VxFS snapshots when the file system is unmounted.</p> <p>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.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>                                                                                                                                                                                                |
| AccessPermissionChk | <p>If the value of this attribute is set to 1 or 2, the monitor verifies that the values of the MntPtPermission, MntPtOwner, and MntPtGroup attributes are the same as the actual mounted file system values. If any of these do not match the values that you have defined, a message is logged.</p> <p>If the value of this attribute is 2, and if the mounted file system permissions do not match the attribute values, the Monitor agent function returns the state as OFFLINE.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p> |

Table 2-14 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CreateMntPt        | <p>If the value of this attribute is 0, no mount point is created. The mount can fail if the mount point does not exist with suitable permissions.</p> <p>If the value of this attribute is 1 or 2, and a mount point does not exist, the agent creates a mount point with system default permissions when the resource is brought online. If the permissions of the mount point is less than 555, a warning message is logged.</p> <p>If the value of this attribute is 2, and the mount point does not exist, the agent creates a mount point with system default permissions when the resource is brought online. If the permissions for the mount point are less than 555, a warning message is logged. In addition, VCS deletes the mount point and any recursively created directories when the resource is brought offline. The mount point gets deleted only if it is empty, which is also true for recursive mount points.</p> <p>Type and dimension: integer-scalar<br/>Default: 0</p> |
| MntPtGroup         | <p>This attribute specifies the group ownership of the mounted file system. The agent verifies the group ownership of the mounted file system every monitor cycle if the value of the AccessPermissionChk attribute is not 0.</p> <p>Type and dimension: string-scalar<br/>Example: "grp1"</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| MntPtOwner         | <p>This attribute specifies the user ownership of the mounted file system. The agent verifies the user ownership of the mounted file system every monitor cycle if the value of the AccessPermissionChk attribute is not 0.</p> <p>Type and dimension: string-scalar<br/>Example: "usr1"</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 2-14** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MntPtPermission    | <p>This attribute specifies the permissions of the mounted file system in absolute format of a four-digit octal. The agent verifies the mode of the mounted file system every monitor cycle if the value of the AccessPermissionChk attribute is not 0.</p> <p>Type and dimension: string-scalar</p> <p>Example: "0755"</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| OptCheck           | <p>The value of this attribute determines if VCS should verify the mount options. The state of the resource is determined based on the result of the verification.</p> <p>If set to 0 (default), the mount options are not checked.</p> <p>If the value of the OptCheck attribute is 1, 2 or 3, a check is performed to see if the mount command options that you have specified for VCS are set in the MountOpt attribute. The MountOpt attributes should be the same as the actual mount command options. If the actual mount options differ from the MountOpt attribute, a message is logged. The state of the resource depends on the value of this attribute.</p> <p>If the value of the attribute is 1, the state of the resource is unaffected.</p> <p>If the value is 2, the state of the resource is set to offline.</p> <p>If the value is 3, state of the resource is set to unknown.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p> |
| RecursiveMnt       | <p>If the value of this attribute is 1, VCS creates all the parent directories of the mount point if necessary.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 2-14 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ReuseMntPt         | <p>If the same mount point needs to be specified in more than one mount resource, set the value of this attribute to 1. Note that this attribute only accepts a value of 1 or 0.</p> <p>To use this attribute, the cluster administrator needs to add this attribute to the arglist of the agent. Set the appropriate group and resource dependencies such that only one resource can come online on a system at a time.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p> |

## Resource type definition

```

type Mount (
 static keylist RegList = { VxFSMountLock }
 static keylist SupportedActions = { "mountpoint.vfd",
 "mounted.vfd", "vxfslic.vfd" , "chgmtlock", "mountentry.vfd" }
 static str ArgList[] = { MountPoint, BlockDevice, FSType,
 MountOpt, FsckOpt, SnapUmount, CkptUmount, SecondLevelMonitor,
 SecondLevelTimeout, OptCheck, CreateMntPt, MntPtPermission,
 MntPtOwner, MntPtGroup, AccessPermissionChk, RecursiveMnt,
 VxFSMountLock }
 static str IMFRegList[] = { MountPoint, BlockDevice, FSType }
 str MountPoint
 str BlockDevice
 str FSType
 str MountOpt
 str FsckOpt
 boolean SnapUmount = 0
 boolean CkptUmount = 1
 boolean SecondLevelMonitor = 0
 int SecondLevelTimeout = 30
 int OptCheck = 0
 int CreateMntPt = 0
 int ReuseMntPt = 0
 str MntPtPermission
 str MntPtOwner
 str MntPtGroup
 int AccessPermissionChk = 0
 boolean RecursiveMnt = 0
 boolean VxFSMountLock = 0
)

```

## Mount agent notes

The Mount agent has the following notes:

- [“High availability fire drill”](#) on page 76
- [“VxFS file system lock”](#) on page 76
- [“TMF usage notes”](#) on page 77
- [“Enabling second level monitoring for the Mount agent”](#) on page 77

### 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/fstab` file

For more information about using the high availability fire drill see the *Veritas Cluster Server Administrator’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
```

## IMF usage notes

If you use IMF for intelligent resource monitoring, review the following recommendations. Depending on the value of the FSType attribute, you must set the MonitorFreq key value of the IMF attribute as follows:

- FSType attribute value is vxfs:
  - For VxFS version 5.1 SP1:  
You can either set the MonitorFreq to 0 or a high value. Setting the value of the MonitorFreq key to a high value will ensure that the agent does not run the monitor function frequently. Setting the MonitorFreq key to 0 will disable the traditional monitoring while IMF monitoring is in progress. Traditional monitoring will be done only after receiving the notification for a resource. However, if the value of the AccessPermissionChk attribute is set to 1, then set the MonitorFreq key value to the frequency at which you want the agent to run the monitor function.
  - For VxFS versions 5.1 5.0.1 or earlier,  
With VxFS versions prior to 5.1 SP1, VCS IMF only monitors file systems getting mounted and unmounted. To monitor other events, you must enable poll-based monitoring. Set the MonitorFreq key value to the frequency at which you want the agent to run the monitor function.
- FSType attribute value is bindfs:  
IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 SP1:
  - IMF should not be enabled for the resources where the BlockDevice can get mounted on multiple MountPoints.
  - If FSType attribute value is nfs, then IMF registration for “nfs” file system type is not supported.

See the *Veritas Cluster Server Administrator's Guide* for the IMF attribute description.

## Enabling second level monitoring for the Mount agent

Second level monitoring can be enabled for the Mount agent only if FSType is set to "nfs".

To enable second level monitoring, run the following commands

```
1 haconf -makerw
2 hares -override resource_name LevelTwoMonitorFreq
3 hares -modify resource_name LevelTwoMonitorFreq 1
```

```
4 haconf -dump -makero
```

For more details about the `LevelTwoMonitorFreq` attribute, refer to the *Veritas Cluster Server Agent Developer's Guide*.

## Sample configurations

### Basic SCSI configuration

```
Mount mnt_r1_1 (
 MountPoint = "/testdir/LVM_R1_1"
 BlockDevice = "/dev/mapper/emc_vg-emc_r1"
 FSType = ext3
 MountOpt = "rw"
 FsckOpt = "-y"
)
```

### VxFS mount lock example

```
Mount test_mnt (
 MountPoint = "/home/export"
 BlockDevice = "/dev/vx/dsk/nfsdg/vol0"
 FSType = vxfs
 MountOpt = rw
 FsckOpt = "-n"
 VxFSMountLock = 1
)
```

### NFS mount example

```
Mount mnt1 (
 MountPoint = "/mnt/vctest"
 BlockDevice = "vcslx597:/root/AMF"
 FSType = nfs
)
```

## Debug log levels

The Mount agent uses the following debug log levels:

DBG\_1, DBG\_2, DBG\_3, DBG\_4, DBG\_5

# Network agents

This chapter contains the following:

- [“About the network agents”](#) on page 79
- [“IP agent”](#) on page 81
- [“NIC agent”](#) on page 87
- [“IPMultiNIC agent”](#) on page 94
- [“MultiNICA agent”](#) on page 102
- [“DNS agent”](#) on page 115

## About the network agents

Use network agents to provide high availability for networking resources. The network agents support IPv4 and IPv6 addresses.

---

**Note:** The `ifconfig` command is deprecated, instead use the `ip` command for network operations. While the `ifconfig` and `ip` both are supported for backward compatibility, Symantec recommends that you use the `ip` command.

---

## Agent comparisons

### IP and NIC agents

The IP and NIC agents:

- Monitor a single NIC

### IPMultiNIC and MultiNICA agents

The IPMultiNIC and MultiNICA agents:

- Operate in two modes:
  - IP Conservation (IPC) Mode, which uses fewer IP addresses
  - Performance Mode (PM), which provides faster failover, but uses more IP addresses
- Monitor single or multiple NICs
- Check the backup NICs at fail over
- Use the original base IP address when failing over
- Have only one active NIC at a time

## 802.1Q trunking

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

The underlying utility to manage 802.1Q trunk interfaces is `vconfig`. For example, you can create a trunk interface on the physical interface:

```
vconfig add eth2 10
```

This creates a trunk interface called `eth2.10` in the default configuration. In this case, the physical NIC `eth2` must be connected to a trunk port on the switch. You can now use `eth2.10` like a regular physical NIC in a NIC, IP, and MultiNICA resource configuration. You can remove it with the following command.

```
vconfig rem eth2.10
```

VCS neither creates nor removes trunk interfaces. The administrator should set up the trunking as per the operating system vendor's documentation rather than using `vconfig` directly.

## 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 Linux bonded interfaces.

## 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
- Checks for the existence of the interface configured in the IP and NIC resources
- Checks for the validity of the configured PrefixLen attribute value for IPv6 configuration

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

## Dependencies

IP resources depend on NIC resources.

**Figure 3-1** Sample service group that includes an IP resource



## Agent functions

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online  | Configures the IP address to the NIC. Checks if another system is using the configured IP address and issues a warning. For IPv4 addresses, it uses the <code>ifconfig</code> command to set the IPv4 address on an unique alias of the interface if the Options attribute is configured and none of the IPOptions and IPRouteOptions attributes are configured, else it makes use of the <code>ip</code> command. For IPv6 addresses, the <code>ip</code> command is used. |
| Action  | The various functions of the action agent are as follows: <ul style="list-style-type: none"><li>■ <code>route.vfd</code><br/>Checks for the existence of a route to the IP from the specified NIC. Checks for the validity of the configured PrefixLen attribute value for IPv6 addresses.</li><li>■ <code>device.vfd</code><br/>Checks for the existence of the interface configured in the Device attribute.</li></ul>                                                    |
| 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.                                                           |
| FAULTED | Indicates that the IP address could not be brought online, usually because the NIC configured in the IP resource is faulted or the IP address was removed out of VCS control. |

## Attributes

**Table 3-1** Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Address            | <p>A virtual IP address that is associated with the interface, and which is different from the base IP address. Note that when the <code>NetMask</code> attribute is configured, the IP address that you specify must not be the same as the configured physical IP address.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"><li>■ IPv4:<br/>"192.203.47.61"</li><li>■ IPv6:<br/>"2001::10"</li></ul> <p><b>Note:</b> For configuring the IP address of a different network than the network of the base IP address, you need to configure the following values in <code>/etc/sysctl.conf</code>.</p> <p>File: <code>/etc/sysctl.conf</code></p> <pre># avoid deleting secondary IPs on deleting the primary IP net.ipv4.conf.default.promote_secondaries = 1 net.ipv4.conf.all.promote_secondaries = 1</pre> <p>After changing these values, load the current values of the <code>/etc/sysctl.conf</code> file using the following command:</p> <pre># sysctl -p /etc/sysctl.conf</pre> <p>This is to make sure that if the primary IP address is unplumbed, the secondary IP address for a network is promoted to the primary address.</p> |
| Device             | <p>The name of the NIC device that is associated with the IP address. Requires the device name without an alias.</p> <p>Type and dimension: string-scalar</p> <p>Example: <code>eth0</code></p> <p>In above example, <code>eth0</code> is specified to assign the IP address to the next available alias of <code>eth0</code>. Use the <code>ip addr</code> command to display a list of NICs that are up and the IP addresses assigned to each NIC.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PrefixLen          | <p>Required to use the IPv6 protocol.</p> <p>See "<a href="#">PrefixLen</a>" on page 84.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 3-2 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NetMask            | <p>The subnet mask that is associated with the IP address. For the IPv4 protocol, specify the value of NetMask attribute in decimal (base 10). If you do not specify the Netmask attribute, the agent uses the operating system's default netmask.</p> <p>Symantec recommends that you configure the value for this attribute. The default value for this attribute can cause delays in establishing inter-node connections.</p> <p>Configure this attribute only if the IP address is an IPv4 address.</p> <p>Type and dimension: string-scalar</p> <p>Example: "255.255.255.0"</p> |
| PrefixLen          | <p>This is the prefix for the IPv6 address represented as the CIDR value.</p> <p>Type-dimension: integer-scalar</p> <p>Range: 0 - 128</p> <p>Default: 1000</p> <p><b>Note:</b> Note that the default value is intentionally invalid for this attribute. You must set the value of this attribute to a range from 0 to 128 to activate this attribute.</p> <p>Example: 64</p>                                                                                                                                                                                                         |
| Options            | <p>Options for the <code>ifconfig</code> command.</p> <p>The <code>ifconfig</code> command generally resembles:</p> <pre>ifconfig dev inet ipv4addr netmask netmask Options up</pre> <p>You must configure either this Options attribute or both the IPOptions and IPRouteOptions attributes.</p> <p>Do not configure any options in this attribute that you can specify in other attributes. For example, for the netmask use the NetMask attribute.</p> <p>Type and dimension: string-scalar</p> <p>Example: "broadcast 172.20.9.255"</p>                                          |

Table 3-2 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPOptions          | <p>Specifies the extra options that are passed to the <code>ip addr add</code> command.</p> <p>The agent uses this attribute in tandem with the <code>IPRouteOptions</code> attribute.</p> <p>The <code>ip addr add</code> command generally resembles:</p> <pre>"ip addr add ipv4addr/prefixlen IPOptions label label dev device"</pre> <p><b>Note:</b> If you configure this attribute, the agent ignores the <code>Options</code> attribute and uses the <code>ip</code> command instead of <code>ifconfig</code>.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ "broadcast 172.20.9.255"</li> <li>■ "scope link"</li> </ul> |
| IPRouteOptions     | <p>Specifies the extra options that are passed to the <code>ip route add</code> command.</p> <p>The agent uses this attribute in tandem with the <code>IPOptions</code> attribute.</p> <p>The <code>ip route add</code> command resembles:</p> <pre>"ip route add IPRouteOptions dev device"</pre> <p><b>Note:</b> If you configure this attribute, the agent ignores the <code>Options</code> attribute and uses the <code>ip</code> command instead of <code>ifconfig</code>.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ "default via 172.20.9.1"</li> <li>■ "scope link"</li> </ul>                                       |

## Resource type definition

```

type IP (
 static keylist SupportedActions = { "device.vfd", "route.vfd" }
 static str ArgList[] = { Device, Address, NetMask, PrefixLen,
 Options, IPOptions, IPRouteOptions }
 str Device
 str Address
 str NetMask
 int PrefixLen = 1000
 str Options

```

```
 str IPOptions
 str IPRouteOptions
)
```

## Sample configurations

### Configuration 1

```
IP IP_192_203_47_61 (
 Device = eth0
 Address = "192.203.47.61"
)
```

### Configuration using specified NetMask

```
IP IP_192_203_47_61 (
 Device = eth0
 Address = "192.203.47.61"
 NetMask = "255.255.248.0"
)
```

## Debug log levels

The IP agent uses the following debug log levels:

DBG\_1, DBG\_2, DBG\_4, DBG\_5

## 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. This resource's Operation value is `None`.

Some NICs maintain their connection status in a hardware register. For NICs that maintain their connection status, the agent uses `ethtool` and `MII` to determine the status of the NIC resource. For NICs that do not maintain their connection status, the agent uses a ping or a broadcast to determine the status of the resource.

For the NIC and IP agents, VCS supports Linux bonded interfaces.

## 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 Administrator's Guide*.

## Dependencies

The NIC resource does not depend on any other resources.

**Figure 3-2** Sample service group that includes a NIC resource



## Bonded network interfaces

The NIC agent now supports using bonded network interfaces.

See “[Monitoring bonded NICs](#)” on page 91.

## Agent functions

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor | <p>If the NIC maintains its connection status, the agent uses MII to determine the status of the resource.</p> <p>If the NIC does not maintain its connection status, the agent verifies that the NIC is configured. The agent then sends a ping to all the hosts that are listed in the NetworkHosts attribute. If the ping test is successful, it marks the NIC resource <code>ONLINE</code>.</p> <p>If the NetworkHosts attribute list is empty, or the ping test fails, the agent counts the number of packets that the NIC received. The agent compares the count with a previously stored value. If the packet count increases, the resource is marked <code>ONLINE</code>. If the count remains unchanged, the agent sends a ping to the broadcast address of the device to generate traffic on the network.</p> <p>The agent counts the number of packets that the NIC receives before and after the broadcast. If the count increases, the resource is marked <code>ONLINE</code>. If the count remains the same or decreases over a period of five broadcast cycles, the resource is marked <code>OFFLINE</code>.</p> <ul style="list-style-type: none"><li>■ <code>device.vfd</code><br/>Checks for the existence of the interface configured in the Device attribute.</li></ul> |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## State definitions

|                      |                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------|
| <code>ONLINE</code>  | Indicates that the NIC resource is working.                                                            |
| <code>FAULTED</code> | Indicates that the NIC has failed.                                                                     |
| <code>UNKNOWN</code> | 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             | <p>Specifies the name of the NIC that you want to monitor.</p> <p>Use the <code>ip addr</code> command to list all network adapters and the IP addresses assigned to each NIC.</p> <p>Type and dimension: string-scalar</p> <p>Example: "eth0" or "eth1"</p> |

**Table 3-4** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mii                | <p>Flag that defines whether the NIC maintains its connection status.</p> <p>If this flag is set to 1, the agent uses ethtool and MII hardware registers, instead of the ping and packet count method. The agent uses this method to determine the health of the network card.</p> <p>If the flag is set to 0, the agent does not use Mii to monitor the status of the NIC.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 1</p> |

**Table 3-4** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NetworkHosts       | <p>List of hosts on the network that receive pings to determine the state of the NIC. Specify the IP address of the host—not the host name.</p> <p>The specified hosts must be pingable:</p> <ul style="list-style-type: none"> <li>■ from all the cluster nodes that are specified in the SystemList attribute for the service group to which the resource belongs</li> <li>■ through the device that is specified in the Device attribute</li> </ul> <p>The command to ping the host (hostip) via a NIC device (nicdev) is:</p> <pre># ping -I nicdev hostip</pre> <p>If more than one network host is listed, the monitor returns ONLINE if the ping test is successful with at least one of the hosts.</p> <p>You can use both IPv4 and IPv6 NetworkHost addresses, and you can configure both types of addresses in the same resource.</p> <p>Type and dimension: string-vector</p> <p>Example:</p> <ul style="list-style-type: none"> <li>■ IPv4:<br/>{ "166.93.2.1", "166.99.1.2" }</li> <li>■ IPv6:<br/>{ "2001::1", "166.93.2.1" }</li> </ul> |
| PingOptimize       | <p>Attribute that defines whether the agent sends a broadcast ping before it retrieves the received packet statistics. This attribute is used when Mii is not set and no network hosts are specified.</p> <p>If the value of this attribute is 1, the agent retrieves received packet statistics from the netstat command and compare them with previously stored values. The agent sends a broadcast ping to the network only if the packet count remains unchanged.</p> <p>If the value of this attribute is 0, the agent sends a broadcast ping before it checks the network statistics.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 1</p>                                                                                                                                                                                                                                                                                                                                                                                            |

## Resource type definition

```
type NIC (
 static keylist SupportedActions = { "device.vfd" }
```

```
static int OfflineMonitorInterval = 60
static str ArgList[] = { Device, PingOptimize, Mii, NetworkHosts
}
static str Operations = None
str Device
int PingOptimize = 1
int Mii = 1
str NetworkHosts[]
)
```

## Monitoring bonded NICs

The NIC agent can monitor the network interfaces (bond0, bond1, etc.) that the bonding driver exports. Refer to operating system vendor documentation to set up the bonds and to configure your system to load the bonding driver correctly.

For monitoring a bond interface, the two important settings are:

- The value of the `miimon` parameter, which you set while loading the bonding driver. `miimon` is a parameter to the bonding module and has a default setting of 0.
- The value of the `Mii` attribute (`Mii`) of the NIC resource, which you set at runtime. `Mii` is an attribute of the NIC resource and has a default setting of 1.

### Setting `Mii` and `miimon`

For the following cases, the name of the monitored bond interface is `B`. If you do not use one of the following cases to set up bonding, the bonding driver can potentially provide incorrect health status. This incorrect health status can result in VCS failing to fault the resource appropriately.

#### Case 1

Accept defaults—`miimon` is 0 and `Mii` is 1. Each of `B`'s slaves must support the `netif_carrier_ok` in-kernel call.

#### Case 2

When you set `miimon` to anything except 0 (`miimon!=0`) and `Mii` to 1, both the hardware and the drivers of each of `B`'s slaves must support the MII-based health monitoring.

#### Case 3

When you set `Mii` to 0, the NIC agent uses ping, which each card supports. In this case, the `miimon` setting is irrelevant.

## Sample configurations

### Configuration for using Mii

If the NIC does not respond to Mii, the agent uses network statistics to monitor the device.

```
NIC groupx_eth0 (
 Device = eth0
 Mii = 1
 PingOptimize = 1
)
```

### Configuration for using network hosts

```
NIC groupx_eth0 (
 Device = eth0
 NetworkHosts = { "166.93.2.1", "166.99.1.2" }
)
```

### IPv6 configuration

The following is a basic configuration for IPv6 with IP and NIC resources. In the following sample, *nic\_value* represents the base NIC value for the platform (for example, eth0).

```
group nic_group (
 SystemList = { sysA = 0, sysB = 1 }
)

NIC nic_resource (
 Device@sysA = eth0
 Device@sysB = eth1
 PingOptimize = 0
 NetworkHosts@sysA = { "2001:db8:c18:2:214:4fff:fe96:11",
 "2001:db8:c18:2:214:4fff:fe96:1" }
 NetworkHosts@sysB = { "2001:db8:c18:2:214:4fff:fe96:1111",
 "2001:db8:c18:2:214:4fff:fe96:111" }
 Protocol = IPv6
)

Phantom phantom_resource (
)

group ip_group (
 SystemList = { sysA = 0, sysB = 1 }
)

IP ip_resource (
 Device@sysA = eth0
 Device@sysB = eth1
 Address = "2001:db8:c18:2:214:4fff:fe96:102"
```

```
 PrefixLen = 64
)

 Proxy proxy_resource (
 TargetResName = nic_resource
)

 ip_resource requires proxy_resource
```

## Debug log levels

The NIC agent uses the following debug log levels:

DBG\_1, DBG\_4, DBG\_5

## 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.

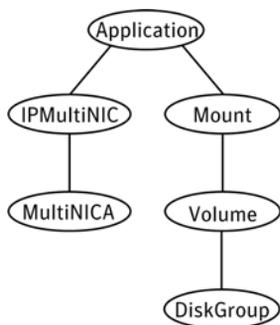
For the IPMultiNIC and MultiNICA agents, VCS supports Linux bonded interfaces.

The IPMultiNIC and MultiNICA agents supports IPv4 and IPv6.

## Dependencies

IPMultiNIC resources depend on MultiNICA resources.

**Figure 3-3** Sample service group that includes 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.<br>Sends out a gratuitous ARP.                                                                               |
| 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.                                                      |
| FAULTED | Indicates that the IP address could not be brought online, usually because all the NICs in the MultiNICA resource are faulted or the IP address was removed out of VCS control. |

## Attributes

**Table 3-5** Required attributes

| Required attribute | Description                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Address            | <p>The virtual IP address that is assigned to the active NIC.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ IPv4: "10.128.10.14"</li> <li>■ IPv6: "2001:DB8::"</li> </ul> |
| MultiNICAResName   | <p>Name of the associated MultiNICA resource that determines the active NIC.</p> <p>Type and dimension: string-scalar</p> <p>Example: "MultiNICA_res1"</p>                                                                         |
| PrefixLen          | <p>Required to use the IPv6 protocol.</p> <p>See "<a href="#">PrefixLen</a>" on page 97.</p>                                                                                                                                       |

**Table 3-6** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NetMask            | <p>For the IPv4 protocol, specify the value of NetMask attribute in decimal (base 10). If you do not specify the Netmask attribute, the agent uses the operating system's default netmask.</p> <p>Symantec recommends that you configure the value for this attribute. The default value for this attribute can cause delays in establishing inter-node connections.</p> <p>Configure this attribute only if the IP address is an IPv4 address.</p> <p>Type and dimension: string-scalar</p> <p>Example: "255.255.255.0"</p> |

Table 3-6 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PrefixLen          | <p>Specifies the prefix for the IPv6 address represented as the CIDR value.</p> <p>When you use the IPv6 protocol, you must configure a value for this attribute.</p> <p>Type-dimension: integer-scalar</p> <p>Range: 0 - 128</p> <p>Default: 1000</p> <p><b>Note:</b> The default value is intentionally invalid for this attribute. You must set the value of this attribute to a range from 0 to 128 to activate this attribute.</p> <p>Example: 64</p>                                                                                                                                                                                                                                                                                 |
| Options            | <p>The <code>ifconfig</code> command options for the virtual IP address.</p> <p>Type and dimension: string-scalar</p> <p>Example: "mtu 2000"</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| IPOptions          | <p>Specifies the extra options that are passed to the <code>ip addr add</code> command. The <code>ip addr add</code> command resembles the following:</p> <ul style="list-style-type: none"> <li>■ IPv4<br/> <pre>"ip addr add ipv4addr/prefixlen IPOptions label label dev device"</pre> </li> <li>■ IPv6<br/> <pre>"ip addr add ipv6addr/prefixlen IPOptions label label dev device"</pre> </li> </ul> <p>Type and dimension: string-scalar</p> <p><b>Note:</b> If you configure this attribute, the agent ignores the Options attribute and uses the <code>ip</code> command instead of <code>ifconfig</code>.</p> <p>Example:</p> <ul style="list-style-type: none"> <li>■ "broadcast 172.20.9.255"</li> <li>■ "scope link"</li> </ul> |

## Resource type definition

```
type IPMultiNIC (
 static int MonitorTimeout = 200
 static int OfflineMonitorInterval = 120
 static int ToleranceLimit = 2
 static str ArgList[] = { Address, NetMask, PrefixLen,
 MultiNICAResName, Options, IPOptions, "MultiNICAResName:Probed"
 }
 str Address
 str MultiNICAResName
 str NetMask
 int PrefixLen = 1000
 str Options
 str IPOptions
)
```

## Sample configuration: IPMultiNIC and MultiNICA

Refer to the MultiNICA agent for more information.

### IPv4 configuration

```
cluster foo (
 UserNames = { admin = "cDRpdxPmHpzS." }
 CounterInterval = 5
)
system sysA (
)
system sysB (
)
group grp1 (
 SystemList = { sysA = 1, sysB = 2 }
)

IPMultiNIC ip1 (
 Address = "192.123.10.177"
 MultiNICAResName = mnic
 NetMask = "255.255.248.0"
)

MultiNICA mnic (
 Device @sysA = { eth0 = "192.123.10.127", eth1 =
 "192.123.11.127" }
 Device @sysB = { eth0 = "192.123.10.128", eth2 =
 "192.123.11.128" }
 NetMask = "255.255.248.0"
 NetworkHosts = { "192.123.10.129", "192.123.10.130" }
)
```

```
ipl requires mnic

// resource dependency tree
//
// group grp1
// {
// IPMultiNIC ip1
// {
// MultiNICA mnic
// }
// }
// }
```

## IPv6 configuration

```
cluster foo (
 UserNames = { admin = "cDRpdxPmHpzS." }
 CounterInterval = 5
)
system sysA (
)
system sysB (
)
group grp1 (
 SystemList = { sysA = 1, sysB = 2 }
)
IPMultiNIC ip1 (
 Address = "2001::110"
 MultiNICAResName = mnic
 PrefixLen=96
)
MultiNICA mnic (
 Device @sysA = { eth0 = "2001::10", eth1 = "2001::11" }
 Device @sysB = { eth0 = "2001::12", eth2 = "2001::13" }
 PrefixLen=96
 NetworkHosts = { "2001::1", "2001::500" }
)
```

```
ipl requires mnic

// resource dependency tree
//
// group grp1
// {
// IPMultiNIC ip1
// {
// MultiNICA mnic
// }
// }
// }
```

## Mixed mode configuration—IPv4 and IPv6

```

cluster foo (
 UserNames = { admin = "cDRpdxPmHpzS." }
 CounterInterval = 5
)
system sysA (
)
system sysB (
)
group grp1 (
 SystemList = { sysA = 1, sysB = 2 }
)

 IPMultiNIC ip1 (
 Address = "2001::110"
 MultiNICAResName = mnic
 PrefixLen=96
)

 IPMultiNIC ip2 (
 Address = "192.123.10.177"
 MultiNICAResName = mnic
 NetMask="255.255.248.0"
)
 MultiNICA mnic (
 Device @sysA = { eth0 = "192.123.10.127", eth1 =
 "192.123.11.127" }
 Device @sysB = { eth0 = "192.123.10.128", eth2 =
 "192.123.11.128" }
 NetMask = "255.255.248.0"
 DualDevice @sysA = { eth0 = "2001::10", eth1 =
 "2001::11" }
 DualDevice @sysB = { eth0 = "2001::12", eth2 =
 "2001::13" }
 PrefixLen=96
 NetworkHosts = { "2001::1", "192.123.10.129" }
)

ip1 requires mnic
ip2 requires mnic

// resource dependency tree
//
// group grp1
// {
// IPMultiNIC ip1
// {
// MultiNICA mnic
// }
// IPMultiNIC ip2
// {
// MultiNICA mnic

```

```
// }
// }
```

## Debug log levels

The IPMultiNIC agent uses the following debug log levels:

DBG\_1, DBG\_4, DBG\_5

## 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.

The IPMultiNIC agent depends upon the MultiNICA agent to select the most preferred NIC on the system. IPMultiNIC brings the virtual IP online or offline. However, if the MultiNICA resource changes its active device, the MultiNICA agent handles the shifting of IP addresses.

If a NIC on a system fails, the MultiNICA agent selects another active NIC. The agent then shifts the virtual IP address to the newly selected active NIC. Only in a case where all the NICs that form a MultiNICA agent fail, does the virtual IP address shift to another system.

If you associate an interface with a MultiNICA resource, do not associate it with any other MultiNICA 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, and
- The Proxy resources that point to the MultiNICA resource in the other service groups.

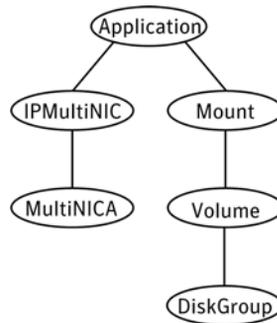
The MultiNICA agent can operate in two modes “[IP Conservation Mode \(ICM\)](#)” on page 103 and “[Performance Mode \(PM\)](#)” on page 103. With sufficient IP addresses, use PM.

The IPMultiNIC and MultiNICA agents support Linux bonds.

The IPMultiNIC and MultiNICA agents support IPv4 and IPv6.

## Dependencies

The MultiNICA resource does not depend on any other resources.

**Figure 3-4** Sample service group that includes a MultiNICA resource

## IP Conservation Mode (ICM)

Requires fewer IP addresses than Performance Mode, but provides slower failover.

### Configuration

When a MultiNICA resource controls all the NICs of a cluster node, the NICs must have the same base IP address. This IP address must be unique, and cannot appear on any other NIC on any other node. You do not need to enable the base IP addresses beforehand. This mode does not support failing back the NIC, see the optional Failback attribute.

### Operation

When you specify all the NICs with the same base IP address, the agent runs in ICM. It enables the base IP address on the active NIC.

In case of a failover, it moves the base IP address to the new active NIC. It also moves all the virtual IP addresses that are configured on that NIC. It then tries to find the first working NIC in the order of priority.

## Performance Mode (PM)

Requires more IP addresses than ICM, but provides faster failover. You do not have to spend time enabling and disabling base IP addresses and reinstating lost routes, thus no resultant service disruption occurs.

### Configuration

The MultiNICA agent controls each NIC, and each NIC must have a unique base IP address. The base IP address cannot appear on any other NIC on the same

node or any other node. The base IP address of all the devices in a single MultiNICA resource must belong to the same subnet in the configuration.

When you configure a single NIC under a MultiNICA resource, the MultiNICA agent uses PM. The base IP addresses have to be enabled on each NIC under MultiNICA control. The addresses need to be enabled before starting VCS and handing over the management of the NICs to the agent.

## Operation

The agent uses this mode when all NICs under the MultiNICA agent have separate base IP addresses specified. The mode requires that you enable the base IP addresses before starting VCS. When a NIC goes down, the agent migrates only virtual IP addresses.

In this mode, you can set the Failback attribute to 1 or 0:

- If you set the Failback attribute to 1, in each monitor cycle the agent checks to see if a preferred NIC is up. If the NIC is up, it selects that NIC as the active NIC and moves the virtual IP addresses to the preferred NIC.
- If you set the Failback attribute to 0, the agent selects a new active NIC only if the current active NIC fails. It selects the new active NIC in the order of priority.

## Agent function

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor | Uses ethtool and Medium Independent Interface (MII) to request the device status. If the hardware does not respond, the agent sends a ping to the hosts that are listed in the NetworkHosts attribute. If the ping test fails, the agent checks for activity on a configured interface by sampling the input packets that are received on that interface. If the agent does not detect activity, it forces activity by sending out a broadcast ping. If the agent does not receive a network reply, it migrates to the most suitable next interface. |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## Attributes

While configuring the MultiNICA resource, in addition to the required attributes specified in [Table 3-7, "Required attributes"](#), you must also configure at least one set of attributes from the following:

- For IPv4 configurations:
  - The Options attribute and the RouteOptions attribute, or
  - The IPv4RouteOptions attribute

- For IPv6 configurations:
  - The IPv6RouteOptions attribute

Table 3-7 Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Device             | <p>List of devices and associated base IP addresses. This attribute must be specified separately for each system in the SystemList. You must specify the devices in the list in the order of priority. The first device that the agent determines as “up” becomes the active device, to which the agent assigns a corresponding IP address.</p> <p>For <a href="#">IP Conservation Mode (ICM)</a>: if all the NICs configured in the Device attribute are down, the MultiNICA agent faults the resource after a 2-3 minute interval. This delay occurs because the MultiNICA agent tests the failed NIC several times before it marks the resource offline. The engine log records messages that provide a detailed description of the failover events. Find the engine log in /var/VRTSvcs/log/engine_A.log.</p> <p>Type and dimension: string-association</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ IPv4           <pre>Device@vcslnx1={ eth1 = "10.212.100.178", eth2 = "10.212.100.179" } Device@vcslnx2={ eth2 = "10.212.100.180", eth3 = "10.212.100.181" }</pre> </li> <li>■ IPv6           <pre>Device@vcslnx1={ eth1 = "1234::5678", eth2 = "1234::5679" } Device@vcslnx2={ eth3 = "1234::5680", eth4 = "1234::5681" }</pre> </li> </ul> <p>You can use IPv4 or IPv6 base addresses in the Device attribute, but make sure all of the base addresses use a common IP version.</p> |
| PrefixLen          | <p>Specifies the length of network prefix that is associated with the base IP address.</p> <p><b>Note:</b> This attribute is required for IPv6 configuration.</p> <p>Type and Dimension: Integer scalar</p> <p>Example: 64</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Table 3-7** Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NetMask            | <p>Specifies the netmask that is associated with the base IP address. The value must be specified in decimal (base 10).</p> <p>Configure this attribute only if the IP address is an IPv4 address.</p> <p>Type and dimension: string-scalar</p> <p>Example: "255.255.252.0"</p> |

**Table 3-8** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DualDevice         | <p>The DualDevice attribute specifies the list of devices and associated IPv6 base addresses.</p> <p>Specify:</p> <ul style="list-style-type: none"> <li>■ this attribute separately for each system in the SystemList.</li> <li>■ the devices in the list in the order of priority.</li> </ul> <p>The first device that the agent determines as “up” becomes the active device, to which the agent assigns a corresponding IP address.</p> <p>NICs in Device and DualDevice attributes should be identical and in the same order.</p> <p>Use the DualDevice attribute only when configuring mixed IPv4/IPv6 stacks. In that case, use the Device attribute to configure the IPv4 stack and the DualDevice attribute to configure the IPv6 stack.</p> <p>Example:</p> <pre>DualDevice@vcslnx1={ eth1 = 2001::DB8, eth2 = 2001::DB9} DualDevice@vcslnx2={ eth3 = 2001::DB10, eth4 = 2001::DB11}</pre> |
| Failback           | <p>This attribute determines if the active NIC should be changed to a preferred NIC, even though the current NIC is healthy. If operating in the ICM mode, change the value to 0.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 3-8 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv4AddrOptions    | <p>The extra options that are passed to the <code>ip addr add</code> command for IPv4 addresses.</p> <p>The agent uses this attribute in tandem with the IPv4RouteOptions attribute.</p> <p>The <code>ip addr add</code> command generally resembles:<br/><code>"ip -4 addr add <i>ipv4addr/prefixlen</i> <i>IPOptions</i> device <i>dev</i>"</code></p> <p><b>Note:</b> If you configure the Options attribute, the agent ignores this attribute.</p> <p>Type and dimension: string-scalar</p> <p>Example:</p> <ul style="list-style-type: none"><li>■ "broadcast 172.20.9.255"</li><li>■ "scope link"</li></ul> |
| IPv4RouteOptions   | <p>The extra options that are passed to the <code>ip route add</code> command for IPv4 addresses.</p> <p>The agent uses this attribute in tandem with the IPv4AddrOptions attribute.</p> <p>The <code>ip route add</code> command generally resembles:<br/><code>"ip route add <i>Route</i> <i>IPRouteOptions</i>"</code></p> <p><b>Note:</b> If you configure the Options attribute, the agent ignores this attribute.</p> <p>Type and dimension: string-scalar</p> <p>Example:</p> <ul style="list-style-type: none"><li>■ "default via 172.20.9.1"</li><li>■ "scope link"</li></ul>                            |

**Table 3-8** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv6AddrOptions    | <p>The extra options that are passed to the <code>ip addr add</code> command for IPv6 addresses. The agent uses this attribute in tandem with the IPv6RouteOptions attribute.</p> <p>The <code>ip addr add</code> command resembles:</p> <pre>"ip -6 addr add ipv6addr/prefixlen IPv6AddrOptions device dev"</pre> <p>Type and dimension: string-scalar</p> <p>Example: "scope link"</p>                                                                                                                                                                                                    |
| IPv6RouteOptions   | <p>The extra options that are passed to the <code>ip route add</code> command for IPv6 addresses. The <code>ip route add</code> command generally resembles:</p> <pre>"ip route add Route IPv6RouteOptions"</pre> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ "default via 2001::DB2"</li> <li>■ "scope link"</li> </ul>                                                                                                                                                                                                             |
| LinkOptions        | <p>Specifies options for the <code>ip link</code> command, which can bring an interface up or down. The <code>ip link</code> command generally resembles:</p> <pre>"ip link dev up LinkOptions"</pre> <p>Type and dimension: string-scalar</p> <p>Example:</p> <ul style="list-style-type: none"> <li>■ "mtu 1024"</li> <li>■ "broadcast 172.20.9.255"</li> </ul>                                                                                                                                                                                                                           |
| NetworkHosts       | <p>List of hosts on the network that receive pings to determine the state of the NICs. Specify the IP address of the host, not the host name. Include the hosts that all the NICs in the Device list can reach. If more than one network host is listed, monitor returns ONLINE if the ping test is successful with at least one of the hosts.</p> <p>You can use both IPv4 and IPv6 NetworkHost addresses, and you can configure both types of addresses in the same resource.</p> <p>Type and dimension: string-vector</p> <p>Example: NetworkHosts = { "2001::1", "192.123.10.129" }</p> |

Table 3-8 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Options            | <p>The <code>ifconfig</code> options that you want to use when you assign the base IP address to the active device.</p> <p>This attribute does not support IPv6.</p> <p><b>Note:</b> If you configure this Options attribute, the agent does not use the <code>IPv4Options</code> and <code>IPv4RouteOptions</code> attributes.</p> <p>Type and dimension: string-scalar</p> <p>Example: "broadcast 10.212.100.255"</p>                                                                                                                                                                                                                                       |
| PingOptimize       | <p>Determines whether or not a broadcast ping is sent before checking network statistics, which are used to determine the state of the NIC (if MII is not supported and the ping to <code>NetworkHosts</code> does not confirm the NIC is up.) A value of 1 indicates a broadcast ping does not occur, a value of 0 indicates a broadcast ping occurs.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 1</p>                                                                                                                                                                                                                                        |
| RouteOptions       | <p>Assignment of a base IP address to a device, which is followed by a <code>route add</code> command. The command has the options specified for this attribute. <code>RouteOptions</code> are applicable only when configuring the local host as the default gateway. No routes are added if this string is set to <code>NULL</code>.</p> <p>The <code>RouteOptions</code> attribute works in tandem with the <code>Options</code> attribute. You must configure the <code>Options</code> attribute when you use this attribute or the agent ignores this attribute.</p> <p>Type and dimension: string-scalar</p> <p>Example: "default gw 166.98.16.103"</p> |

## Resource type definition

```

type MultiNICA (
 static int MonitorTimeout = 240
 static str ArgList[] = { Device, DualDevice, NetMask, PrefixLen,
 Options, RouteOptions, PingOptimize, MonitorOnly, NetworkHosts,
 Failback, LinkOptions, IPv4AddrOptions, IPv6AddrOptions,
 IPv4RouteOptions, IPv6RouteOptions, Mii }
 static str Operations = None
 str Device{}
 str DualDevice{}
 str NetMask
 int PrefixLen = 1000
 str Options
 str RouteOptions
 str LinkOptions
 str IPv4AddrOptions
 str IPv6AddrOptions
 str IPv4RouteOptions
 str IPv6RouteOptions
 int PingOptimize = 1
 str NetworkHosts[]
 boolean Failback = 1
 boolean Mii = 1
)

```

## Sample configurations

### MultiNICA and IPMultiNIC Performance Mode configuration

In this example, two systems (sysA and sysB) each have a pair of network interfaces (eth0 and eth1, eth0 and eth2). These interfaces have different physical IP addresses and the agent behaves in [Performance Mode \(PM\)](#).

The MultiNICA resource fails over only the logical IP address to the backup NIC in the event of a failure. The resource ip1 has the Address attribute, which contains the logical IP address. In the event of a NIC failure on sysA, the logical IP address fails over from eth0 to eth1. In the event that eth1 fails—the address fails back to eth0—as long as eth0 is reconnected.

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 fails over to sysB.

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

```

cluster foo (

```

```
UserNames = { admin = "cDRpdxPmHpzS." }
CounterInterval = 5
)

system sysA (
)
system sysB (
)

group grp1 (
 SystemList = { sysA = 1, sysB = 2 }
)

IPMultiNIC ip1 (
 Address = "192.123.10.177"
 MultiNICAResName = mnic
 NetMask = "255.255.248.0"
)

MultiNICA mnic (
 Device @sysA = { eth0 = "192.123.10.127", eth1 =
 "192.123.11.127" }
 Device @sysB = { eth0 = "192.123.10.128", eth2 =
 "192.123.11.128" }
 NetMask = "255.255.248.0"
 NetworkHosts = { "192.123.10.129", "192.123.10.130" }
)

ip1 requires mnic

// resource dependency tree
//
// group grp1
// {
// IPMultiNIC ip1
// {
// MultiNICA mnic
// }
// }
```

## MultiNICA and IPMultiNIC IP Conservation Mode Configuration

In this example, two systems (sysA and sysB) each have a pair of network interfaces (eth0 and eth1, eth0 and eth2). These interfaces have a common physical IP address and the agent behaves in [IP Conservation Mode \(ICM\)](#).

The MultiNICA resource fails over both the physical IP and the logical IP addresses to the backup NIC in the event of a failure. The resource ip1 has the Address attribute, which contains the logical IP address. In the event of a NIC

failure on sysA, the IP addresses fail over from eth0 to eth1. In the event that eth1 fails—the addresses fail back to eth0—if eth0 is reconnected.

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 fails over to sysB.

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

```
cluster foo (
 UserNames = { admin = "cDRpdxPmHpzS." }
 CounterInterval = 5
)

system sysA (
)
system sysB (
)

group grp1 (
 SystemList = { sysA = 1, sysB = 2 }
)
IPMultiNIC ip1 (
 Address = "192.123.10.177"
 MultiNICAResName = mnic
 NetMask = "255.255.248.0"
)

MultiNICA mnic (
 Device @sysA = { eth0 = "192.123.10.127", eth1 =
"192.123.10.127" }
 Device @sysB = { eth0 = "192.123.10.128", eth2 =
"192.123.10.128" }
 NetMask = "255.255.248.0"
 NetworkHosts = { "192.123.10.129", "192.123.10.130" }
 Failback = 0
)

ip1 requires mnic

// resource dependency tree
//
// group grp1
// {
// IPMultiNIC ip1
// {
// MultiNICA mnic
```

```
// }
// }
```

## IPv6 configuration

The following sample is for IPv6 use.

```
group mnica_group (
 SystemList = { sysA = 0, sysB = 1 }
)

IPMultiNIC ipmnic_res (
 Address = "2007:192::1627:161"
 MultiNICAResName = mnica_res
 PrefixLen = 64
)

MultiNICA mnica_res (
 Device @sysA = { eth0 = "fe80::214:4fff:fe96:ae0a",
 eth1 = "fe80::214:4fff:fe96:ae0a" }
 Device @sysB = { eth0 = "fe80::214:4fff:fe96:ae0b",
 eth2 = "fe80::214:4fff:fe96:ae0b" }
 PrefixLen = 64
)

ipmnic_res requires mnica_res
```

## Mixed mode configuration—IPv4 and IPv6

```
cluster foo (
 UserNames = { admin = "cDRpdxPmHpzS." }
 CounterInterval = 5
)

system sysA (
)

system sysB (
)

group grp1 (
 SystemList = { sysA = 1, sysB = 2 }
)
IPMultiNIC ip1 (
 Address = "2001::110"
 MultiNICAResName = mnic
 PrefixLen=96
)
IPMultiNIC ip2 (
 Address = "192.123.10.177"
 MultiNICAResName = mnic
 NetMask="255.255.248.0"
```

```
)
MultiNICA mnic (
 Device @sysA = { eth0 = "192.123.10.127", eth1 =
 "192.123.11.127" }
 Device @sysB = { eth0 = "192.123.10.128", eth1 =
 "192.123.11.128" }
 NetMask = "255.255.248.0"
 DualDevice @sysA = { eth0 = "2001::10", eth1 = "2001::10" }
 DualDevice @sysB = { eth0 = "2001::11", eth1 = "2001::11" }
 PrefixLen=96
 NetworkHosts = { "2001::1", "192.123.10.129" }
)

ip1 requires mnic
ip2 requires mnic
// resource dependency tree
//
// group grp1
// {
// IPMultiNIC ip1
// {
// MultiNICA mnic
// }
// IPMultiNIC ip2
// {
// MultiNICA mnic
// }
// }
// }
```

## Debug log levels

The MultiNICA agent uses the following debug log levels:  
DBG\_1, DBG\_2, DBG\_3, DBG\_4, DBG\_5

## 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, 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.

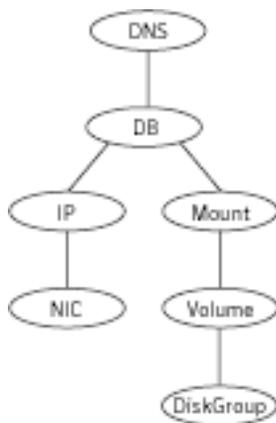
For important information about this agent, refer to:

“[DNS agent notes](#)” on page 122

## Dependencies

No dependencies exist for the DNS resource.

**Figure 3-5** Sample service group that includes a DNS resource



## Agent functions

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online  | <p>Updates one or more name servers with the resource records.</p> <p>The agent updates the name servers defined in the StealthMasters attribute. If you have not configured this attribute then the agent obtains the name of the master server by sending an Start of Authority (SOA) query. This query retrieves the SOA record of the zone defined in the agent's Domain attribute. This SOA record contains the name of the master server.</p> <p>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 server serves the forward (A or AAAA) and reverse zones.</p> <p>Finally the agent generates an Online lock file to indicate that the resource is online on the current system.</p> |
| Offline | <p>Removes the Online lock file.</p> <p>If attribute OffDelRR is true, offline removes all records that the ResRecord keys define.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Monitor | <p>Returns the ONLINE state if at least one name server reports all mappings that ResRecord defines. The name servers are the master or StealthMaster servers and all the servers for which an NS record for the zone exists.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Clean   | <p>Removes the Online lock file, if it exists.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Open    | <p>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 records that the ResRecord attribute defines. If all the name servers fail to report all the records, the agent function removes the Online lock file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Action  | <p>Different action agent functions follow:</p> <ul style="list-style-type: none"><li>■ keyfile.vfd<br/>This action entry point checks if the key file as specified in the TSIGKeyFile attribute exists either locally or on shared storage.</li><li>■ dig.vfd<br/>This action entry point checks if dig and nsupdate binaries exist and are executable.</li><li>■ master.vfd<br/>This action entry point checks if stealth masters are able to reply to SOA query for the configured domain.</li></ul>                                                                                                                                                                                                                                                                                                                   |

## State definitions

|         |                                                                                                                                                                                                                                  |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ONLINE  | Online lock file exists and at least one name server can return all configured resource records.                                                                                                                                 |
| OFFLINE | Indicates an offline state when at least one of the following is true: <ul style="list-style-type: none"><li>■ The online lock does not exist.</li><li>■ None of the name servers can report all of the RRs' mappings.</li></ul> |
| 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 3-9** Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Domain             | <p>A string representing the DNS zone that the agent administers. The domain name can only contain alphanumeric symbols and the dash.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"><li>■ Forward mapping<br/>"demo.example.com"</li><li>■ IPv4 reverse mapping<br/>"2.168.192.in-addr.arpa"</li></ul> |

Table 3-9 Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ResRecord          | <p>ResRecord is an association of DNS resource record values. Each ResRecord attribute consists of two values: <i>DNS record key</i> = <i>DNS record data</i>. Note that the record key must be a unique value.</p> <p>If the resource record list contains any invalid value as a part of the record key or a record data of the ResRecord attribute, the resource reports an UNKNOWN state.</p> <p>Type and dimension: string-association</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ For forward mapping, where the zone is demo.example.com: <ul style="list-style-type: none"> <li>- sles901 = "192.168.2.191"</li> <li>- ww2 = sles901</li> <li>- sles9ip6 = "2007::1:2:3:abc"</li> </ul> </li> <li>■ For a multi-home DNS record, typically for one host with two network interfaces and different addresses, but the same DNS name. This results in two-A records, or a single A record with continuation lines. <ul style="list-style-type: none"> <li>sle902 = "192.168.2.102 10.87.13.22"</li> </ul> A multi-home AAAA DNS record can be configured as follows: <ul style="list-style-type: none"> <li>sle902 = "1234::5678 1234::AABB:CCDD"</li> </ul> </li> <li>■ For reverse IPv4 address mapping, where the zone is 2.168.192.in-addr.arpa: <ul style="list-style-type: none"> <li>191 = "sles901.demo.example.com."</li> </ul> </li> <li>■ For reverse IPv6 address mapping, where the zone is 3.0.0.0.2.0.0.0.1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.7.0.0.2.ip6.arpa: <ul style="list-style-type: none"> <li>cba = "sles9ip6.demo.example.com."</li> </ul> </li> </ul> <p>Use only partial host names. If you use a fully qualified domain name, append a period "." at the end of the name.</p> <p>For CNAME records, use:</p> <ul style="list-style-type: none"> <li>■ ResRecord = { www = mydesktop }<br/>or</li> <li>■ ResRecord = { www = "mydesktop.marketing.example.com." }<br/>Where the Domain attribute is "marketing.example.com"</li> </ul> |

**Table 3-10** Required attributes

| Required attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ResRecord (cont.)  | <p>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:</p> <ul style="list-style-type: none"> <li>■ PTR: if the Domain attribute ends with .arpa</li> <li>■ A: if the record data field is four sets of numbers, where a period separates each set. The following details the pattern it tries to match: [1-223].[0-255].[0-255].[0-255] Hexadecimal is not supported.</li> <li>■ AAAA: if the record data fields are in multiple sets of hexadecimal format, then this record is an IPv6 associated type AAAA record.</li> <li>■ CNAME: for any other valid record data.</li> </ul> <p><b>Note:</b> 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. As an exception to this, the DNS agent allows underscore character ("_") in hostnames. Make sure that the DNS server supports the underscore character before you configure any DNS resource records to have the underscore character in their hostnames.</p> |

**Table 3-11** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TTL                | <p>A non-zero integer represents the “Time To Live” value, in seconds, for the DNS entries in the zone that you want to update.</p> <p>A lower value means more hits on your DNS server, while a higher value means more time for your clients to learn about changes.</p> <p>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.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 86400</p> <p>Example: 3600</p> |

**Table 3-11** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| StealthMasters     | <p>The list of primary master name servers in the domain.</p> <p>This attribute is optional since the first name server is retrieved from the zone's SOA (Start of Authority) record.</p> <p>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.</p> <p>Type and dimension: string-vector</p> <p>Example: { "10.190.112.23" }</p> |
| TSIGKeyFile        | <p>Required when you configure DNS for secure updates. Specifies the absolute path to the file containing the private TSIG (Transaction Signature) key.</p> <p>Type and dimension: string-scalar</p> <p>Example:</p> <pre>/var/tsig/example.com.+157+00000.private</pre>                                                                                                                                                                                                                                                                 |
| CreatePTR          | <p>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, make sure that the same master or stealth servers must serve the forward (A or AAAA) and reverse zones.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>                                                                                                         |

**Table 3-11** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OffDelRR           | <p>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.</p> <p>The online agent function always adds records if they do not exist.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p> |

## Resource type definition

```

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

```

## DNS agent notes

The DNS agent has the following notes:

- [“High availability fire drill”](#) on page 123
- [“Monitor scenarios”](#) on page 123
- [“Sample Web server configuration”](#) on page 123
- [“Secure DNS update for BIND 9”](#) on page 124
- [“Setting up secure updates using TSIG keys for BIND 9”](#) on page 124

## 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 can respond to the SOA query made 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 Administrator'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 3-12** 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.example.com` that maps to the canonical name `location1.example.com`. The browser retrieves the IP address for the web server by querying a domain name server. If the web server fails over from location one to location two (`location2.example.com`), the domain name servers need a new canonical name mapping for `www.example.com`. The `www.example.com` alias is now updated to point to the canonical name of the standby system in location two.

## 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 -b 128 -n ZONE veritas.com.
```

- 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"
)
```

## Sample configurations

This section contains sample configurations for this agent.

### Basic IPv6 configuration

This sample configuration provides basic configuration for IPv6 support. In the following sample, `nic_value` represents the base NIC value for the platform (for example, `en0`, `bge0`, `eth0`, etc.)

```
group ipv6_group_dns (
 SystemList = { sysA = 0, sysB = 1 }
)

DNS ipv6group_dns_res (
 Critical = 0
 Domain = "ipv6.vcs.net"
 TSIGKeyFile =
 "/var/tsig/Kipv6.vcscd.net.+157+18435.private"
 StealthMasters = { "2001:db8:c18:2:69c4:3251:bac1:6cbe" }
 ResRecord = {
 vcssystemCv6 = "2001:db8:c18:2:214:4fff:fe96:8833",
 sysC = vcssystemCv6 }
)

IP ipv6group_ip_res (
 Device @sysA = nic_value
 Device @sysB = nic_value
 Address = "2001:db8:c18:2:214:4fff:fe96:8833"
 PrefixLen = 64
)

NIC ipv6group_nic_res (
 Device @sysA = nic_value
 Device @sysB = nic_value
 NetworkHosts = { "2001:db8:c18:2:214:4fff:fea2:fd50" }
)

ipv6group_dns_res requires ipv6group_ip_res
ipv6group_ip_res requires ipv6group_nic_res
```

## IPv6 CNAME sample configuration

The following sample configuration uses CNAME values.

```
group cname_group (
 SystemList = { sysA = 0, sysB = 1 }
)

DNS cname_group_dns_res (
 Domain = "example.com"
 StealthMasters = { "3ffe:556::1000:5761" }
 ResRecord @sysA = { ftp = foo }
 ResRecord @sysB = { ftp = bar }
 OffDelRR = 1
)
```

## IPv4 A sample configuration

The following sample configuration uses A values.

```
group forwardv4_group (
 SystemList = { sysA = 0, sysB = 1 }
)

DNS forward_group_v4_resource (
 Domain = "example.com"
 StealthMasters = { "3ffe:556::1000:5761" }
 ResRecord @sysA = { www = "10.200.56.240" }
 ResRecord @sysB = { www = "10.200.56.244" }
 OffDelRR = 1
)
```

## Debug log levels

The DNS agent uses the following debug log levels:

DBG\_1, DBG\_2, DBG\_3, DBG\_4, DBG\_5

# File share agents

This chapter contains the following:

- [“About the file service agents”](#) on page 127
- [“NFS agent”](#) on page 128
- [“NFSRestart agent”](#) on page 133
- [“Share agent”](#) on page 140
- [“About the Samba agents”](#) on page 145
- [“SambaServer agent”](#) on page 148
- [“SambaShare agent”](#) on page 153
- [“NetBios agent”](#) on page 156

## 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`, `mountd`, `statd`, and `lockd` daemons required by all exported NFS file systems. Configure the NFS resource in a separate parallel service group with the `AutoStart` attribute set to 1.

You should configure only a single 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 must use a Proxy resource. The Proxy resource can point to the NFS resource in the first group. Duplicate NFS resources will cause a problem when the NFS resources are brought online concurrently—only the NFS resource started first will be successfully brought online, while the rest of the NFS resources may report online failure.

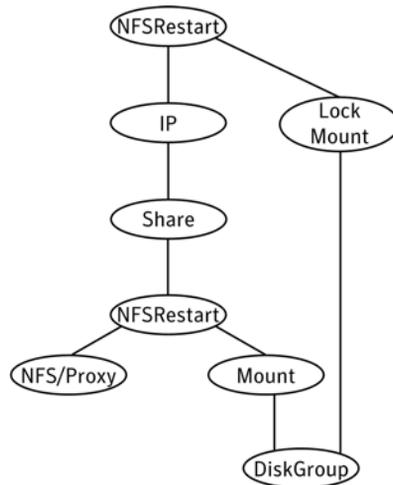
For important information about this agent, refer to:

“[NFS agent notes](#)” on page 131

## Dependencies

For more information regarding NFS resource dependencies, refer to the *Veritas Cluster Server Administrator’s Guide*.

**Figure 4-1** Sample service group that includes an NFS resource



## Agent functions

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online  | <p>Starts nfsd and mountd daemons for all kernels. If daemons are not running, the agent starts the daemons and exits.</p> <ul style="list-style-type: none"><li>■ For Red Hat, the agent also starts the lockd and statd daemons.</li><li>■ For SUSE, the agent also starts the lockd daemon. The sm-notify daemon sends lock recovery notifications, which then perform their jobs and terminate automatically.</li></ul>                                                                                                                            |
| Monitor | <p>Monitors versions 2 and 3 of the nfsd daemon and versions 1, 2, and 3 of the mountd daemon for all kernels.</p> <ul style="list-style-type: none"><li>■ For Red Hat, the agent also monitors the lockd and statd daemons. Monitors versions 1, 3, 4 of the lockd daemon and version 1 of the statd daemon. When the value of the NFSv4Support attribute is 1, nfsd version 4 is also monitored.</li><li>■ For SUSE, the agent also monitors the lockd daemon. Monitors versions 1, 3, 4 of the lockd daemon. Monitors version 1 of statd.</li></ul> |
| Clean   | <p>Stops nfsd and mountd daemons for all kernels.</p> <ul style="list-style-type: none"><li>■ For Red Hat, the agent also stops the lockd and statd daemons.</li><li>■ For SUSE, the agent also stops the lockd daemon.</li></ul>                                                                                                                                                                                                                                                                                                                      |

## 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 attributes | Description                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GracePeriod         | <p>Required when the value of the NFSLockFailover attribute for NFSRestart resource is 1. GracePeriod specifies the amount of time, in seconds, that lock recovery is allowed by the NFS server after its reboot.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 90</p>                                                                                                           |
| LockFileTimeout     | <p>The NFS and the NFSRestart agents require a synchronization mechanism when the group to which they belong is in transition, for example going online or coming offline. A file serves as this synchronization mechanism. The LockFileTimeout attribute specifies the maximum time that the synchronization file exists.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 180</p> |
| Nproc               | <p>Specifies the number of concurrent NFS requests that the server can handle.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 8</p> <p>Example: 16</p>                                                                                                                                                                                                                            |
| NFSSecurity         | <p>Specifies whether to start the NFS security daemon rpc.svcgssd or not.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>                                                                                                                                                                                                                                                    |
| NFSv4Support        | <p>Specifies whether to start the NFSv4 daemon rpc.idmapd or not and whether to monitor nfsd version 4.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>                                                                                                                                                                                                                      |

## Resource type definition

```
type NFS (
 static int RestartLimit = 1
 static str Operations = OnOnly
 static str ArgList[] = { Nproc, GracePeriod, NFSSecurity,
 NFSv4Support, LockFileTimeout }
 int Nproc = 8
 int GracePeriod = 90
 boolean NFSSecurity = 0
 boolean NFSv4Support = 0
 int LockFileTimeout = 180
)
```

## NFS agent notes

The NFS agent has the following notes:

- [“Prerequisites for NFS lock recovery”](#) on page 131
- [“Using NFSv4”](#) on page 131

### Prerequisites for NFS lock recovery

If you plan on using lock recovery on a Linux system, store locking information on shared storage so that it is accessible to the system where NFS fails over. Using this information, NFS carries out lock recovery.

For more information, refer to the NFSRestart agent.

### Using NFSv4

The NFS agent provides NFSv4 support to export shares using the attribute NFSv4Support. The shared directory that has the fsid=0 option becomes the root of all exports. The client needs to mount only this root file system instead of mounting all shares individually.

The syntax is:

```
mount -t nfs4 <server>:/ <mountpoint>
```

Always use a slash (/) to end the path after the colon (:).

All the file systems, other than the root file system, needs to have the nohide option set in Options attribute of share resources. Set the nohide option so that authentic clients can seamlessly move through the tree of exported file systems by mounting the root file system.

To enable NFSv4 support on your node, you must have the rpc\_pipefs (pipe file system) mounted on the node. At boot time, rpc\_pipefs is mounted on every Linux node. In situations where it is not mounted, mount rpc\_pipefs on the cluster node.

### To mount `rpc_pipefs`

- ◆ At the prompt on the node, enter the following:

```
mount -t rpc_pipefs rpc_pipefs /var/lib/nfs/rpc_pipefs
```

## Sample configurations

On each node in your cluster, you can find sample NFS, NFSRestart, and Share configurations in `/etc/VRTSvcs/conf/sample_nfs/`.

For more information regarding agent configuration, refer to the *Veritas Cluster Server Administrator's Guide*.

## Debug log levels

The NFS agent uses the following debug log levels:

DBG\_1, DBG\_2, DBG\_3, DBG\_4, DBG\_5

# NFSRestart agent

The NFSRestart agent provides the following functionalities:

- Manages NFS lock recovery service by recovering the NFS record locks after sudden server crash.
- Prevents potential NFS ACK storms by terminating NFS server services before offline of NFS VIP to close all TCP connections with the NFS client.

If you have configured the NFSRestart agent for lock recovery, the NFSRestart agent starts the smsyncd daemon. The daemon copies the NFS locks from the local directory `/var/lib/nfs` to shared storage. The agent's online function copies the locks from shared storage to local directory `/var/lib/nfs`.

For important information about this agent, refer to:

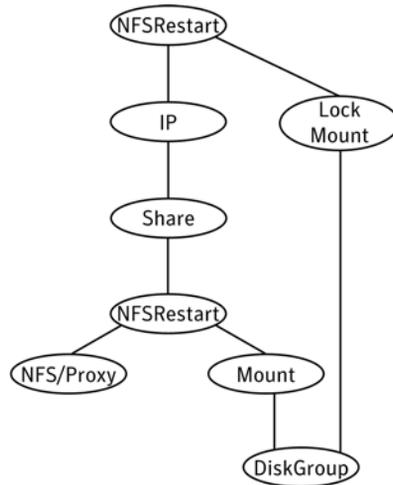
[“NFSRestart agent notes”](#) on page 137

## Dependencies

For more information regarding NFSRestart resource dependencies, refer to the *Veritas Cluster Server Administrator's Guide*.

You must use two NFSRestart resources in a service group. The lower NFSRestart resource must have its Lower attribute set to 1. The upper NFSRestart resource should be at the top of the resource dependency tree and the lower NFSRestart resource should be below the Share resource in the resource dependency tree. The NFSRestart resources and the Share resources must be inside the same service group.

Figure 4-2 Sample service group that includes an NFSRestart resource



## Agent functions

- Online
- For the lower NFSRestart resource:
- If the value of the NFSLockFailover attribute is 1, the agent terminates statd and lockd.
- For the upper NFSRestart resource:
- If the value of the NFSLockFailover attribute is 1, the agent copies the NFS record locks from shared storage to /var/lib/nfs directory.
  - Starts the statd and lockd daemons.
  - Starts the smsyncd daemon to copy the contents of /var/lib/nfs directory to the shared storage (LocksPathName) at regular two second intervals.
  -
- Monitor
- For the lower NFSRestart resource:
- The monitor agent function does nothing.
- For the upper NFSRestart resource:
- If the value of the NFSLockFailover attribute is 1, the agent monitors smsyncd daemon. It restarts the smsyncd daemon if it is not running.
  -

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Offline | <p>For the lower NFSRestart resource:</p> <ul style="list-style-type: none"><li>■ Restarts all the NFS daemons that the upper NFSRestart resource stopped previously.</li></ul> <p>For the upper NFSRestart resource:</p> <ul style="list-style-type: none"><li>■ Terminates the statd and lockd daemons to clear the lock state.</li><li>■ Terminates the nfsd and mountd daemons to close the TCP/IP connections.</li><li>■ Terminates the smsyncd daemon if the daemon is running.</li></ul> |
| Clean   | <p>For the lower NFSRestart resource:</p> <ul style="list-style-type: none"><li>■ Restarts all the NFS daemons that the upper NFSRestart resource stopped previously.</li></ul> <p>For the upper NFSRestart resource:</p> <ul style="list-style-type: none"><li>■ Terminates the statd and lockd daemons to clear the lock state.</li><li>■ Terminates the nfsd and mountd daemons to close the TCP/IP connections.</li><li>■ Terminates the smsyncd daemon if the daemon is running.</li></ul> |
| Action  | <ul style="list-style-type: none"><li>■ nfsconf.vfd<br/>Checks the runlevel information of the system service nfslock to confirm that the lock daemons do not come online automatically after reboot.</li><li>■ lockdir.vfd<br/>Verifies that the NFS lock directory (which is specified by the LocksPathName attribute of NFSRestart) is on shared storage.</li></ul>                                                                                                                          |

## 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** Required attributes

| Required attribute | Description                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NFSRes             | Name of the NFS resource on the system. This attribute is required if the value of the NFSLockFailover attribute is 1.<br>Type and dimension: string-scalar<br>Example: "nfsres1" |

**Table 4-3** 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.<br>Type and dimension: string-scalar<br>Example: "/share1x" |
| NFSLockFailover    | A flag that specifies whether the user wants NFS locks to be recovered after a failover.<br>Type and dimension: boolean-scalar<br>Default: 0                                                                                                                                                                                                                                                               |
| Lower              | Defines the position of NFSRestart resource in the service group. The NFSRestart resource below the Share resource needs a value of 1.<br>The NFSRestart resource on the top of the resource dependency tree has a Lower attribute value of 0.<br>Type and dimension: integer-scalar<br>Default: 0                                                                                                         |

## Resource type definition

```
type NFSRestart (
 static keylist SupportedActions = { "lockdir.vfd", "nfsconf.vfd"
 }
 static str ArgList[] = { "NFSRes:Nproc", "NFSRes:GracePeriod",
 "NFSRes:NFSv4Support", NFSLockFailover, LocksPathName, Lower,
 State }
 str NFSRes
 int Lower = 0
 str LocksPathName
 boolean NFSLockFailover = 0
)
```

## NFSRestart agent notes

The NFSRestart agent has the following notes:

- [“About high availability fire drill”](#) on page 137
- [“Providing a fully qualified host name”](#) on page 137

### About 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 Administrator’s Guide*.

### Providing a fully qualified host name

You must provide a fully qualified host name, for example, nfserver.example.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 (nfserver), NFS lock recovery may fail.

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 `/var/statmon/sm` directory on the NFS client should also contain a fully qualified domain name of the NFS server after the acquisition of locks. Otherwise you need to stop and start the status daemon and lock daemon 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

On each node in your cluster, you can find sample NFS, NFSRestart, and Share configurations in `/etc/VRTSvcs/conf/sample_nfs/`.

For more information regarding agent configuration, refer to the *Veritas Cluster Server Administrator's Guide*.

### Basic agent configurations

For NFS lock recovery:

```
NFSRestart nfsrestart (
 NFSRes = nfsres
 LocksPathName="/shared_mnt/lockinfo"
 NFSLockFailover = 1
 Lower = 0
)

NFSRestart nfsrestart_L (
```

```
NFSRes = nfsres
LocksPathName="/shared_mnt/lockinfo"
NFSLockFailover = 1
Lower = 1
)
For no NFS lock recovery:
NFSRestart nfsrestart (
 NFSRes = nfsres
)

NFSRestart nfsrestart_L (
 NFSRes = nfsres
 Lower = 1
)
```

## Debug log levels

The NFSRestart agent uses the following debug log levels:

DBG\_1, DBG\_3, DBG\_4, DBG\_5

## 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 exported are on shared disks.

For important information on this agent, refer to:

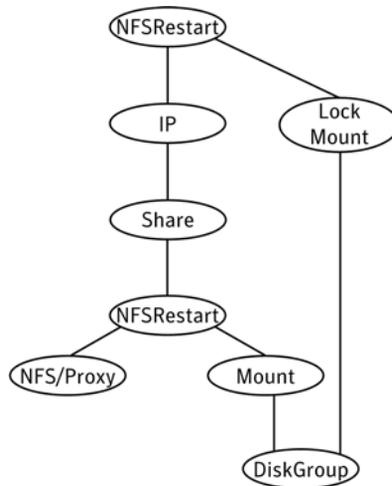
“[Share agent notes](#)” on page 143

## Dependencies

For more information regarding Share resource dependencies, refer to the *Veritas Cluster Server Administrator's Guide*.

Share resources depend on NFS. In an NFS service group, the IP family of resources depends on Share resources.

**Figure 4-3** Sample service group that include 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  | <code>direxists.vfd</code><br>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 4-4** Required attributes

| Required attribute | Description                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------|
| PathName           | Pathname of the file system to be shared.<br>Type and dimension: string-scalar<br>Example: "/share1x" |
| NFSRes             | This attribute has been deprecated.                                                                   |

**Table 4-5** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client             | The Share agent accepts as many clients as the user wishes provided all the clients are exported the same 'PathName'.<br>Client or host where the directory specified by PathName is exported. The client can be a wild card (*) or a fully qualified domain name (FQDN) including the host name.<br>Type and dimension: string-scalar<br>Example: If "outland" is the host name, the FQDN hostname is outland.symantec.com. |
| Options            | Options to the <code>exportfs</code> command. When specifying multiple options, separate them with commas, for example:<br>"rw, no_root_squash"<br>For more information about the <code>exportfs</code> command and its options, refer to the <code>exportfs</code> manpage.<br>Type and dimension: string-vector<br>Default = "ro, async, wdelay, root_squash"                                                              |

**Table 4-5** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OtherClients       | <p>The Client attribute can be assigned one FQDN host name, whereas multiple FQDN host names can be assigned to the 'OtherClients' field.</p> <p>A combination of 'Client' and 'OtherClients' can be used to specify the host names.</p> <p>If both of the Client and OtherClients attributes are left unspecified, the PathName is exported to the world (*).</p> <p>Type and dimension: string-vector</p> |

## Resource type definition

```

type Share (
 static keylist SupportedActions = { "direxists.vfd" }
 static str ArgList[] = { PathName, Client, OtherClients,
 Options, "NFSRes:State" }
 str PathName
 str Client
 str OtherClients[]
 str Options
 str NFSRes
)

```

## 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 Administrator's Guide*.

## Sample configurations

On each node in your cluster, you can find sample NFS, NFSRestart, and Share configurations in `/etc/VRTSvcs/conf/sample_nfs/`.

For more information regarding agent configuration, refer to the *Veritas Cluster Server Administrator's Guide*.

## Debug log levels

The Share agent uses the following debug log levels:

DBG\_1, DBG\_3, DBG\_4, DBG\_5

## 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.
- The default path of the `smbd` and `nmbd` daemons is:  
`/usr/sbin`  
For more information on configuring these paths, refer to the description of the `SambaTopDir` attribute.
- 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

VCS Samba suite of agents support Samba version 3.0 and above. Please check your samba version using the following command:

```
smbctl -V
```

VCS supports most versions of Samba that are bundled with supported operating systems. For operating systems that do not come bundled with Samba, VCS supports most versions that are compatible with the operating system.

## Notes for configuring the Samba agents

The following notes describe configuration considerations for the Samba agents.

### Enabling VCS to detect services started and stopped by smb

Edit the samba file to enable VCS to detect the Samba services that are started and stopped using the following commands:

- `/etc/init.d/smb start` OR  
`service smb start`
- `/etc/init.d/smb stop` OR  
`service smb stop`

Edit the `/etc/sysconfig/samba` file. Add `-s smb.conf-filepath` to `SMBDOPTIONS`. This change is required. The SambaServer agent monitors the resource by matching the configuration file present in running process's arguments list with the value configured in `ConfFile` attribute.

For example:

```
$ cat /etc/sysconfig/samba
Options to smbd
SMBDOPTIONS="-D -s /etc/samba/smb.conf"
Options to nmbd
NMBDOPTIONS="-D -s /etc/samba/smb.conf"
Options for winbindd
WINBINDOPTIONS=""
```

### Configuring multiple SambaServer resources

For configuring multiple SambaServer resources, configure the `SocketAddress` attribute with the unique value of the address where the respective samba daemon listens for connections. Configure the SambaServer resource as a parent resource of the IP resource. Configure this IP resource with the `SocketAddress` attribute value.

## Configuring Samba for non-standard configuration files or non-standard lock directories

Configure the PidFile attribute if you use a non-standard configuration file for Samba or if the lock directory (the directory where Samba pid file resides) for Samba is different than the default location. Use the following command to check the standard locations for the Samba configuration file and the lock directory:

### To check for the default value of the Samba configuration file

- ◆ Enter the following command:

```
smbd -b | grep CONFIGFILE
```

### To check for the default location of the Samba pidfile

- ◆ Enter the following command:

```
smbd -b | grep PIDDIR
```

## SambaServer agent

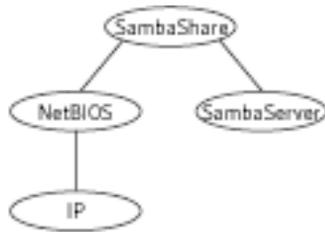
The SambaServer agent starts, stops, and monitors the `smbd` process as a daemon. You can use the agent to make a `smbd` daemon highly available.

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 that includes a SambaServer resource



### Agent functions

|         |                                                                                                                                                                                                                                          |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online  | Starts the <code>smbd</code> daemon at specified or default ports.                                                                                                                                                                       |
| Offline | Stops the <code>smbd</code> daemon.                                                                                                                                                                                                      |
| Monitor | Verifies that the <code>smbd</code> daemon is running by reading its <code>pid</code> 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 <code>smbd</code> 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.                                                                                                                                                                                                   |
| FAULTED | Indicates that the smbd daemon has stopped unexpectedly or is not responding (if in-depth monitoring is enabled) outside of VCS control.                                                                                                                                  |

## Attributes

**Table 4-6** Required attributes

| Required attribute | Description                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ConfFile           | Complete path of the configuration file that Samba uses.<br>Type and dimension: string-scalar<br>Example: "/etc/samba/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.<br>Type and dimension: string-scalar<br>Example:<br>"/var/run" |

**Table 4-7** Optional attributes

| Optional attribute        | Description                                                                                                                                                                                                                                                                                                                            |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IndepthMonitorCyclePeriod | Number of monitor cycles after which the in-depth monitoring is performed. For example, the value 5 indicates that the agent monitors the resource in-depth every five monitor cycles. The value 0 indicates that the agent will not perform in-depth monitoring for the resource.<br>Type and dimension: integer-scalar<br>Default: 5 |
| Ports                     | Ports where Samba accepts connections.<br>To run Samba over NBT (NetBios over TCP/IP), set this attribute to 139. To run Samba directly over TCP/IP, set this attribute to 445.<br>Type and dimension: integer-vector<br>Default: 139, 445                                                                                             |

Table 4-7 Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ResponseTimeout    | <p>Number of seconds the agent waits to receive the session response packet after sending the session request packet. For example, the value 5 indicates that the agent waits for five seconds before receiving the session response packet. Configure this attribute if in-depth monitoring is enabled.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 10</p>                                                                                                     |
| SambaTopDir        | <p>Parent path of Samba daemon and binaries.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/usr"</p>                                                                                                                                                                                                                                                                                                                                                              |
| PidFile            | <p>The absolute path to the Samba daemon pid file. This file contains the process ID of the monitored smbd process.</p> <p>Configure this attribute if you are using a non-standard configuration file name or path. If this agent is not configured for non-standard configuration file names, the agent checks the <code>smbd-ConfFile.pid</code> file for monitoring the resource.</p> <p>Type and dimension: string-scalar</p> <p>Example:</p> <p>"/var/run/smbd.pid"</p> |
| SocketAddress      | <p>The IP address where the Samba daemon (smbd) listens for connections.</p> <p>Configure the SocketAddress attribute if you are configuring multiple SambaServer resources on a node.</p> <p><b>Note:</b> Only IPv4 addresses are supported.</p> <p>Type and Dimension: string-scalar</p> <p>Example: "10.128.10.14"</p>                                                                                                                                                     |

## Resource type definitions

```
type SambaServer (
 static str ArgList[] = { ConfFile, LockDir, Ports,
 IndepthMonitorCyclePeriod, ResponseTimeout, SambaTopDir,
 PidFile, SocketAddress}
 str ConfFile
 str LockDir
 int Ports[] = { 139, 445 }
 int IndepthMonitorCyclePeriod = 5
 int ResponseTimeout = 10
 str SambaTopDir
 str PidFile
 str SocketAddress
)
```

## Sample configurations

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

## Debug log levels

The SambaServer agent uses the following debug log levels:  
DBG\_1, DBG\_2, DBG\_3, DBG\_4, DBG\_5

## SambaShare agent

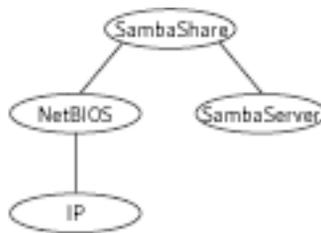
The SambaShare agent adds, removes, and monitors a share by modifying the specified Samba configuration file. You can use the agent to make a Samba Share highly available.

Each filesystem 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 the 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 <code>smbclient</code> 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. |
| FAULTED | Indicates that the share has become unavailable outside of VCS control.               |
| UNKNOWN | Indicates that the agent could not determine the state of the resource.               |

## Attributes

**Table 4-8** Required attributes

| Required attribute | Description                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SambaServerRes     | Name of the SambaServer resource.<br>Type and dimension: string-scalar<br>Example: "smb_res1"                                                                                                                                           |
| ShareName          | Name of the share resource.<br>Type and dimension: string-scalar<br>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 (;).<br>Type and dimension: string-scalar<br>Example: "path=/shared; public=yes; writable=yes" |

## Resource type definition

```

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

```

## Sample configuration

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

## Debug log levels

The SambaShare agent uses the following debug log levels:  
DBG\_1, DBG\_5

## NetBios agent

The NetBios agent starts, stops, and monitors the nmbd daemon. You can use the agent to make the nmbd daemon highly available.

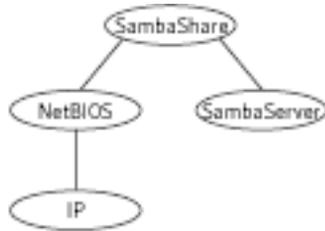
The agent sets, monitors, and resets the names and network interfaces by which the Samba server is known. The agent also sets, monitors and resets Samba to act as a WINS server or domain master or both.

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

## Dependencies

The NetBios resource depends on the IP or the IPMultiNIC resource if the virtual IP address configured in the IP/IPMultiNIC resource is being used in the Interfaces attribute of the NetBios resource.

**Figure 4-6** Sample service group that includes 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: <ul style="list-style-type: none"><li>■ NetBIOS name is not advertised.</li><li>■ A NetBIOS alias is not advertised.</li><li>■ Samba is not handling requests on one of the specified interfaces.</li><li>■ If WINS support is configured, Samba is not providing WINS service.</li><li>■ If domain support is set, Samba is not providing Domain Master service.</li></ul> |
| UNKNOWN | Indicates that the agent could not determine the state of the resource.                                                                                                                                                                                                                                                                                                                                             |
| FAULTED | Indicates that the nmbd daemon has stopped unexpectedly outside of VCS control.                                                                                                                                                                                                                                                                                                                                     |

## Attributes

**Table 4-9** Required attributes

| Required attribute | Description                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------|
| NetBiosName        | Name by which the Samba server is known in the network.<br>Type and dimension: string-scalar  |
| SambaServerRes     | Name of the SambaServer resource.<br>Type and dimension: string-scalar<br>Example: "smb_res1" |

**Table 4-10** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interfaces         | List of network interfaces on which Samba handles browsing.<br>Type and dimension: string-vector<br>Example: "172.29.9.24/16"<br><br><b>Note:</b> Note that if you have configured the SocketAddress attribute value for the corresponding SambaServer resource, then you must also configure the same value paired with the appropriate netmask in the list of interfaces. |
| NetBiosAliases     | List of additional names by which the Samba server is known in the network.<br>Type and dimension: string-vector<br>Example: "host1_samba, myname"                                                                                                                                                                                                                          |
| WinsSupport        | If set to 1, this flag causes the agent to configure Samba as a WINS server.<br>Type and dimension: integer-scalar<br>Default: 0                                                                                                                                                                                                                                            |

**Table 4-10** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DomainMaster       | <p>If the value of this attribute is 1, the agent sets Samba as Domain Master.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p>                                                                                                                                                                                                                                                                                                                            |
| PidFile            | <p>The absolute path to the NetBIOS daemon pid file. This file contains the process ID of the monitored nmbd process.</p> <p>Configure this attribute if you are using a non-standard configuration file name or path. If this agent is not configured for non-standard configuration file names, the agent checks for the nmbd-<i>ConfFile</i>.pid file for resource monitoring.</p> <p>Type and dimension: string-scalar</p> <p>Example:</p> <p>"/var/run/nmbd.pid"</p> |

## Resource type definition

```

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

```

## Sample configuration

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

## Debug log levels

The NetBios agent uses the following debug log levels:  
DBG\_1, DBG\_5

# Service and application agents

This chapter contains the following agents:

- [“About the service and application agents”](#) on page 161
- [“Apache Web server agent”](#) on page 162
- [“Application agent”](#) on page 176
- [“CoordPoint agent”](#) on page 186
- [“Process agent”](#) on page 190
- [“ProcessOnOnly agent”](#) on page 195

## 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.

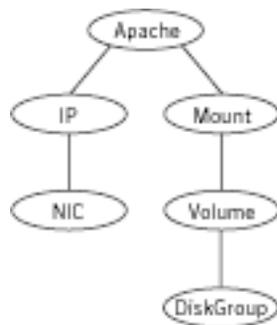
---

For more information regarding this agent:  
See [“Apache Web server notes”](#) on page 169.

## 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 | <p>To stop the Apache HTTP server, the agent:</p> <ul style="list-style-type: none"><li>■ Executes the httpdDir/httpd program with the appropriate arguments (Apache v2.0), or</li><li>■ Sends a TERM signal to the HTTP Server parent process (Apache v1.3).</li></ul> <p>When you specify a file with the EnvFile attribute, the file is sourced before the agent executes the httpd command.</p> |
| 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  | <p>checkconffile.vfd</p> <p>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.</p>                                                                                     |

## State definitions

|         |                                                                                                                                                                                                                                  |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ONLINE  | Indicates that the Apache server is running.                                                                                                                                                                                     |
| OFFLINE | <p>Indicates that the Apache server is not running.</p> <p>Can also indicate that the administrator has stopped the Web server gracefully. Note that the agent uses the PidFile attribute for intentional offline detection.</p> |
| 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.<br>Type and dimension: string-scalar<br>Example: "/apache/server1/conf/httpd.conf"                                                                                                                                                                                                                                                                                                                                                   |
| httpdDir           | Full path of the directory to the httpd binary file<br>Type and dimension: string-scalar<br>Example: "/apache/server1/bin"                                                                                                                                                                                                                                                                                                                                                                                         |
| ResLogLevel        | Controls the agent's logging detail for a specific instance of a resource. Values are: <ul style="list-style-type: none"><li>■ ERROR: Logs error messages.</li><li>■ WARN: Logs error and warning messages.</li><li>■ INFO: Logs error, warning, and informational messages.</li><li>■ TRACE: Logs error, warning, informational, and trace messages. Trace logging is verbose. Use for initial configuration or troubleshooting.</li></ul> Type and dimension: string-scalar<br>Default: INFO<br>Example: "TRACE" |
| PidFile            | This attribute is required when you want to enable the detection of a graceful shutdown outside of VCS control.<br>See " <a href="#">PidFile</a> " on page 167.                                                                                                                                                                                                                                                                                                                                                    |

**Table 5-2** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DirectiveAfter     | <p>A list of directives that httpd processes after reading the configuration file.</p> <p>Type and dimension: string-association</p> <p>Example: DirectiveAfter{} = { KeepAlive=On }</p>                                                                                                                                                                                                                                                                                                                |
| DirectiveBefore    | <p>A list of directives that httpd processes before it reads the configuration file.</p> <p>Type and dimension: string-association</p> <p>Example: DirectiveBefore{} = { User=nobody, Group=nobody }</p>                                                                                                                                                                                                                                                                                                |
| User               | <p>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.</p> <p>Type and dimension: string-scalar</p> <p>Example: "apache1"</p>                                                                                                                                                                                                                                                                                        |
| EnableSSL          | <p>Set to 1 (true) to have the online agent function add support for SSL by including the option <code>-DSSL</code> in the start command. For example:<br/><code>/usr/sbin/httpd -f path_to_httpd.conf -k start -DSSL</code></p> <p>Where <code>path_to_httpd.conf</code> file is the path to the <code>httpd.conf</code> file.</p> <p>Set to 0 (false) it excludes the <code>-DSSL</code> option from the command.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: "1"</p> |

**Table 5-2** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HostName           | <p>The virtual host name that is assigned to the Apache server instance. The host name is used in second-level monitoring for benchmarking the Apache HTTP server.</p> <p>You can use IPv4 or IPv6 addresses for the HostName attribute.</p> <p><b>Note:</b> The HostName attribute is only required when the value of SecondLevelMonitor is 1 (true).</p> <p>Type and dimension: string-scalar</p> <p>Example: "web1.example.com"</p>                                                    |
| Port               | <p>Port number where the Apache HTTP server instance listens. The port number is used in second-level monitoring for benchmarking the Apache HTTP server. Specify this attribute only if SecondLevelMonitor is set to 1 (true).</p> <p>Type and dimension: integer-scalar</p> <p>Default: 80</p> <p>Example: "80"</p>                                                                                                                                                                     |
| EnvFile            | <p>Full path and file name of the file that is sourced before executing httpdDir/httpd. With Apache 2.0, the file <i>ServerRoot/bin/envvars</i>, 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.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/apache/server1/bin/envvars"</p> |

**Table 5-2** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PidFile            | <p>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.</p> <p>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).</p> <p>Type and dimension: string-scalar</p> <p>Example: <code>/var/run/httpd.pid</code></p> |
| SharedObjDir       | <p>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 <code>-R</code> option when executing the httpd program. Refer to the httpd man pages for more information about the <code>-R</code> option.</p> <p>Type and dimension: boolean-scalar</p> <p>Example: <code>"/apache/server1/libexec"</code></p>                                        |
| SecondLevelMonitor | <p>Enables second-level monitoring for the resource. Second-level monitoring is a deeper, more thorough state check of the Apache HTTP server. Valid attribute values are 1 (true) and 0 (false). Specifying this attribute is required.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: <code>"1"</code></p>                                                                                                                                                                                                                     |

**Table 5-2** Optional attributes

| Optional attribute | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SecondLevelTimeout | <p>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.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 30</p> |

**Table 5-3** Resource type attribute

| Required attribute | Description                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| IntentionalOffline | For information on how to use the IntentionalOffline resource type attribute, refer to the <i>Veritas Cluster Server Administrator's Guide</i> . |

## Resource type definition

```

type Apache (
 static keylist SupportedActions = { "checkconf.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 boolean 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”](#) on page 169
- [“About detecting application failure”](#) on page 170
- [“About bringing an Apache Web server online outside of VCS control”](#) on page 170
- [“About high Availability fire drill”](#) on page 170

## 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 170.
- 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.
- 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.

## About 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 the Apache agent uses the Apache Benchmarking utility "ab" to perform detail monitoring. If the exit code of the "ab" utility is 0 and if the command output contains "Benchmarking *HostName*", the agent considers the server online, else 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 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 Administrator's Guide*.

## Sample configurations

### Running two versions of httpd for Linux

This example shows how two versions of `httpd` can run from different locations. In group `Apache_1`, `httpd` runs from Port 80, the default location. The configuration file in `/usr/local/apache/conf/httpd.conf` should indicate `DocumentRoot`, `address`, `port`, and other parameters. In group `Apache_2`, `httpd` runs from `/home/web/apache`. The PID file for this is created in `/home/web/apache/logs/httpd.pid`. The configuration file in `/home/web/apache/conf/httpd.conf` should define parameters for running this version of `httpd`.

Each Apache resource requires an online IP resource. In this example, each Apache resource requires an online mount resource to mount block devices from disks reserved by the Disk Reservation agent.

```
system sysa
```

```
system sysb
```

```
group Apache_1 (
 SystemList = { sysa ,sysb}
 AutoStartList = { sysa}
)

Apache myapacheWeb (
 httpdDir = "/mnt/apache/bin"
 SecondLevelMonitor = 1
 ConfigFile = "/mnt/apache/conf/httpd.conf"
 HostName = "server1.example.com"
 Port = 80
)

IP myapacheIP(
 Device = "eth0"
 Address="192.168.50.50"
 NetMask="255.255.255.0"
)

NIC myapacheNIC (
 Device="eth0"
 NetworkHosts={"172.29.9.178", "172.29.9.179"}
)
```

```
Mount myapacheMnt (
 MountPoint="/mnt/apache/"
 BlockDevice="/dev/sdd2"
 FSType = ext3
 FsckOpt = "-y"
)

DiskReservation myapacheDiskRes(
 Disks = { "/dev/sdd" }
)

myapacheMnt requires myapacheDiskRes
myapacheIP requires myapacheNIC
myapacheWeb requires myapacheIP
myapacheWeb requires myapacheMnt

group Apache_2 (
 SystemList = { sysa,sysb}
 AutoStartList = { sysa}
)

Apache myapacheWeb2(
 httpdDir = "/mnt/apache1/bin"
 SecondLevelMonitor = 1
 ConfigFile = "/mnt/apache1/conf/httpd.conf"
 HostName = "server2.example.com"
 Port = 8080
)

IP myapacheIP2(
 Device = "eth1"
 Address="192.168.60.50"
 NetMask="255.255.255.0"
)

NIC myapacheNIC2(
 Device="eth1"
)

Mount myapacheMnt2(
 MountPoint="/mnt/apache1/"
 BlockDevice="/dev/sdc3"
 FSType = ext3
 FsckOpt = "-y"
)

DiskReservation myapacheDiskRes2(
 Disks = { "/dev/sdc" }
)

myapacheMnt2 requires myapacheDiskRes2
myapacheIP2 requires myapacheNIC2
```

```
myapacheWeb2 requires myapacheIP2
myapacheWeb2 requires myapacheMnt2
```

### Sample main.cf file

```
include "types.cf"

cluster Cluster1 (
 UserNames = { admin = xxxxxx }
)

system SystemA (
)
system SystemB (
)

group Web1 (
 SystemList = { SystemA = 0, SystemB = 1 }
)

 DiskGroup Web1_dg (
 DiskGroup = web1
)

 Volume Web1_vol (
 DiskGroup = web1
 Volume = volweb1
)

 IP Web1_ip (
 Device = eth0
 Address = "10.212.88.220"
 NetMask = "255.255.254.0"
)

 Mount Web1_mnt (
 MountPoint = "/apache/srvr01"
 BlockDevice = "/dev/vx/dsk/web1/volweb1"
 FSType = vxfs
 FckOpt = "-y"
)

 NIC Web1_nic (
 Device = eth0
)

 Apache Web1_http (
 HostName = spartan
 Port = 80
 SecondLevelMonitor = 1
 SecondLevelTimeout = 25
 httpdDir = "/apache/srvr01/bin"
```

```
 EnvFile = "/apache/srvr01/bin/envvars"
 PidFile = /apache/srvr01/log/httpd.pid"
 ConfigFile = "/apache/srvr01/conf/httpd.conf"
 IntentionalOffline = 1
)

Web1_ip requires Web1_nic
Web1_mnt requires Web1_vol
Web1_vol requires Web1_dg
Web1_http requires Web1_ip
Web1_http requires Web1_mnt
```

## Basic IPv6 configuration

The following is a basic IPv6 configuration for the resource.

```
group ipv6group (
 SystemList = { sysA = 0, sysB = 1 }
)

Apache ipv6group_apache_res (
 HostName = "fd4b:454e:205a:110:211:25ff:fe7e:118"
 PidFile = "/myapache/apache/logs/httpd.pid"
 httpdDir = "/myapache/apache/bin"
 ConfigFile = "/myapache/apache/conf/httpd.conf"
 ResLogLevel = TRACE
 SecondLevelTimeout = 20
 IntentionalOffline = 1
)

DiskGroup ipv6group_dg_res (
 DiskGroup = dg01
)

IP ipv6group_ip_res (
 Device = eth0
 Address = "fd4b:454e:205a:110:211:25ff:fe7e:118"
 PrefixLen = 64
)

Mount ipv6group_mnt_res (
 MountOpt = rw
 FsckOpt = "-n"
 BlockDevice = "/dev/vx/dsk/dg01/vol01"
 MountPoint = "/myapache/apache"
 FSType = vxfs
)

NIC ipv6group_nic_res (
 Device = eth0
)

Volume ipv6group_vol_res (
 Volume = vol01
 DiskGroup = dg01
)

ipv6group_apache_res requires ipv6group_mnt_res
ipv6group_apache_res requires ipv6group_ip_res
ipv6group_mnt_res requires ipv6group_vol_res
ipv6group_vol_res requires ipv6group_dg_res
ipv6group_ip_res requires ipv6group_nic_res
```

## 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 bundled, enterprise, or 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

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

## 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 (program.vfd)
- Execution permissions for the specified program (program.vfd)
- The existence of the specified user on the host (user.vfd)
- The existence of the same binary on all nodes (cksum.vfd)

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

## Dependencies

Depending on how you plan to use it, an Application type of resource can depend on IP and Mount resources. Alternatively, instead of the IP resource you can also use the IPMultiNIC resource.

**Figure 5-2** Sample service group that includes an Application resource



## Agent functions

- Online** Runs the command or script that you specify in the value of the StartProgram attribute. Runs the command with the specified parameters in the context of the specified user.
- To bring the resource online, the agent function performs the command:
- ```
su [-] user -c command_to_online_resource
```
- Offline** Runs the command or script that you specify in the value of the StopProgram attribute. Runs the command with the specified parameters in the context of the specified user.
- To take the resource offline, the agent function performs the command:
- ```
su [-] user -c command_to_offline_resource
```
- 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 of 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.
- If the MonitorProgram attribute is specified to monitor the resource, the agent function performs the command:
- ```
su [-] user -c command_to_monitor_resource
```
- imf_init** Initializes the agent to interface with the asynchronous monitoring framework (AMF) kernel driver. This function runs when the agent starts up.
- imf_getnotification** Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel driver. The agent continuously waits for notification and takes action on the resource upon notification.
- imf_register** Registers the resource entities, which the agent must monitor, with the AMF kernel driver. For example, the function registers the PID for online monitoring of a process. This function runs for each resource after the resource goes into steady state (online or offline). The Application agent uses IMF for the processes configured with PidFiles and the MonitorProcesses attribute.

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`.

To forcefully stop the resource, the agent function performs the command:

```
su [-] user -c command_to_monitor_resource
```

Note that the agent uses the `su -` option only when the attribute `UseSUDash` is enabled (1). The `UseSUDash` attribute is disabled (0) by default.

State definitions

ONLINE	Indicates that all processes that are specified in the <code>PidFiles</code> and the <code>MonitorProcesses</code> attribute are running and that the <code>MonitorProgram</code> returns ONLINE.
OFFLINE	Indicates that at least one process that is specified in the <code>PidFiles</code> attribute or <code>MonitorProcesses</code> is not running, or that the <code>MonitorProgram</code> returns OFFLINE.
UNKNOWN	Indicates an indeterminable application state or invalid configuration or that the required attributes have not been configured.
FAULTED	Indicates that the process has terminated unexpectedly or <code>MonitorProgram</code> returns “offline” unexpectedly.

Attributes

Table 5-4 Required attributes

Required attribute	Description
StartProgram	<p>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.</p> <p>Note: Do not use the opening and closing ({}) brace symbols in this string.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/usr/sbin/sample_app start"</p>
StopProgram	<p>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.</p> <p>Note: Do not use the opening and closing ({}) brace symbols in this string.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/usr/sbin/sample_app stop"</p>
<p>At least one of the following attributes:</p> <ul style="list-style-type: none"> ■ MonitorProcesses ■ MonitorProgram ■ PidFiles 	<p>See “Optional attributes” on page 181.</p>

Table 5-5 Optional attributes

Optional attribute	Description
CleanProgram	<p>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.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/usr/sbin/sample_app stop"</p>
MonitorProcesses	<p>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.</p> <p>The process name must be the name that the <code>ps -ef</code> command displays for the process.</p> <p>Type and dimension: string-vector</p> <p>Example: "nmbd"</p>
MonitorProgram	<p>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.</p> <p>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.</p> <p>Note: Do not use the opening and closing ({}) brace symbols in this string.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/usr/sbin/sample_app_monitor all"</p>

Table 5-5 Optional attributes

Optional attribute	Description
PidFiles	<p>A list of files that contain the PID (process ID) 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.</p> <p>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.</p> <p>Type and dimension: string-vector</p>
User	<p>The user name 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.</p> <p>Type and dimension: string-scalar</p> <p>Default: root</p> <p>Example: user1</p>
EnvFile	<p>The environment file that should get sourced before running any of the StartProgram, StopProgram, MonitorProgram or CleanProgram.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: /home/username/envfile</p>

Table 5-5 Optional attributes

Optional attribute	Description
UseSUDash	<p>When the value of this attribute is 0, the agent performs an <code>su User</code> command before it executes the <code>StartProgram</code>, the <code>StopProgram</code>, the <code>MonitorProgram</code>, or the <code>CleanProgram</code> agent functions.</p> <p>When the value of this attribute is 1, the agent performs an <code>su - User</code> command before it executes the <code>StartProgram</code>, the <code>StopProgram</code>, the <code>MonitorProgram</code> or the <code>CleanProgram</code> agent functions.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>

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,
    EnvFile, UseSUDash }
  static str IMFRegList[] = { MonitorProcesses, User, PidFiles,
    MonitorProgram }
  str User = root
  str StartProgram
  str StopProgram
  str CleanProgram
  str MonitorProgram
  str PidFiles[]
  str MonitorProcesses[]
  str EnvFile
  boolean UseSUDash = 0
)

```

Application agent notes

Using Application agent with IMF

- Intelligent monitoring is supported for the Application agent only under specific configurations. The complete list of such configurations is provided in [Table 5-6](#).

Table 5-6 Relation of monitoring attributes with IMF modes

MonitorProgram	MonitorProcesses	PidFiles	IMF Monitoring Mode
Not Configured	Not Configured	Not Configured	Not Applicable
Not Configured	Not Configured	Configured	Online Only
Not Configured	Configured	Not Configured	Online, Offline
Not Configured	Configured	Configured	Online, Offline
Configured	Not Configured	Not Configured	No IMF monitoring
Configured	Not Configured	Configured	No IMF monitoring
Configured	Configured	Not Configured	No IMF monitoring
Configured	Configured	Configured	No IMF monitoring

- When multiple processes are configured under the MonitorProcesses attribute and only some of them are running, offline registration with IMF will fail repeatedly until RegisterRetryLimit is reached. In such a scenario, IMF will not be able to determine when the resource goes ONLINE and the agent will monitor the resource in the traditional way.

Sample configurations

Configuration 1

In this example, you configure the executable `sample_app` 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 `app.pid` specifies and the process `sample_app`.

```
Application samba_app (
    User = "root"
    StartProgram = "/usr/sbin/sample_app start"
    StopProgram = "/usr/sbin/sample_app stop"
    PidFiles = { "/var/lock/sample_app/app.pid" }
```

```
        MonitorProcesses = { "sample_app" }  
    )
```

Configuration 2

In this example, since no user is specified, it uses the root user. The executable `sample_app` starts and stops the application using `start` and `stop` as the command line arguments. The executable `sample_app_monitor` monitors the application and uses `all` as its command line argument. The agent also monitors the `sample_app1` and `sample_app2` processes.

```
Application samba_app2 (  
    StartProgram = "/usr/sbin/sample_app start"  
    StopProgram = "/usr/sbin/sample_app stop"  
    CleanProgram = "/usr/sbin/sample_app force stop"  
    MonitorProgram = "/usr/local/bin/sample_app_monitor all"  
    MonitorProcesses = { "sample_app1", "sample_app2" }  
)
```

Debug log levels

The Application agent uses the following debug log levels:

DBG_1, DBG_2, DBG_3, DBG_4, DBG_5

CoordPoint agent

Use the Coordination Point (CoordPoint) agent to monitor the registrations on the different coordination points on each node. You use this agent to provide server-based I/O fencing. The CoordPoint agent is a monitor-only agent that runs on each node within the client cluster. When you have configured a CP server as a coordination point, the CoordPoint agent performs the following tasks:

- Confirms that the CP server coordination point can communicate with the client cluster.
- Validates the node registrations in the CP server database using the `cpsadm` command.

In case the coordination point is a SCSI-3 based disk, the CoordPoint agent uses the `vxfenadm` command to confirm that the registered keys on the disk are intact. The Monitor agent function contains the monitoring functionality for SCSI-3 disks and CP servers.

If the agent detects an anomaly, the agent reports it to you so you can repair the coordination point. You may have to perform an online coordinator point replacement procedure if the problem is isolated to the keys registered.

Note: The CoordPoint agent that runs on a given client cluster node monitors the keys for coordination points visible to that node alone.

For important information about this agent, refer to:

[“Notes for the CoordPoint agent”](#) on page 188

Dependencies

No dependencies exist for the CoordPoint resource.

Agent functions

Monitor Enables the CoordPoint agent to validate the node registrations in the coordination points and confirms that the coordination points are accessible. CoordPoint resources are persistent, which means that they cannot be brought online or taken offline. They can only monitor the coordination point registrations. For this reason, the service group that contains the CoordPoint resource appears to be offline after a command such as `hastatus -sum`.

The CoordPoint agent also performs I/O fencing reporting activities.

See [“CoordPoint agent I/O fencing reporting activities”](#) on page 188.

State definitions

ONLINE	Indicates that the CoordPoint resource is working.
UNKNOWN	Indicates the agent cannot determine the coordination points resource's state. This state may be due to an incorrect configuration.
FAULTED	Indicates that the number of coordination points with missing keys (or registrations) has exceeded the value of the FaultTolerance attribute.

Attributes

Table 5-7 Required attributes

Required attribute	Description
FaultTolerance	<p>The FaultTolerance attribute determines when the CoordPoint agent declares that the registrations on the coordination points are missing.</p> <p>If the number of coordination points with missing keys (or registrations) exceeds the value of the FaultTolerance attribute, then the agent reports <code>FAULTED</code>.</p> <p>Set the value of this attribute depending on your own configuration requirements. For example, if the FaultTolerance value is set to 1, then the CoordPoint agent reports <code>FAULTED</code> if it sees 2 or more number of coordinator points with missing keys (or registrations).</p> <p>Change the value of the FaultTolerance attribute either before the CoordPoint agent starts to monitor or while the CoordPoint agent is monitoring. If the attribute is set while the CoordPoint agent is monitoring, then the CoordPoint agent reads the new value in the next monitor cycle.</p> <p>To view the current FaultTolerance value, enter the following command:</p> <pre># hares -display coordpoint-res -attribute FaultTolerance</pre> <p>Type and dimension: integer-scalar Default: "0"</p>

Resource type definition

```
type CoordPoint (  
    static str ArgList[] = { FaultTolerance }  
    static int InfoInterval = 300  
    static int OfflineMonitorInterval = 60  
    static str Operations = None  
    int FaultTolerance  
)
```

Notes for the CoordPoint agent

CoordPoint agent I/O fencing reporting activities

The CoordPoint agent also performs the following I/O fencing reporting activities:

- Checks to determine if I/O fencing is running.
If I/O fencing is not running, then the CoordPoint agent reports failure.
- Checks the mode of fencing operation. I/O fencing can operate in one of the following three modes:
 - SCSI-3 mode: If I/O fencing runs in SCSI-3 mode, then the CoordPoint agent continues to monitor.
 - Customized mode: If I/O fencing runs in Customized Fencing mode, then the CoordPoint agent continues to monitor.
 - Disabled mode: If I/O fencing runs in disabled mode, no action is required. The CoordPoint agent returns success.

AutoStartList attribute

AutoStartList is a service group attribute that needs to be populated with a system list. The VCS engine brings up the specified service group on the nodes in the list.

AutoStartList is not a required attribute for the service group that contains the CoordPoint resource. The CoordPoint resource is a persistent resource and when a service group is configured with this type of resource, it cannot be brought online.

Specifying the AutoStartList with a system list does not change the behavior of the service group. The service group will be reflected in OFFLINE status itself, irrespective of the AutoStartList attribute.

Sample configuration

In this example, the coordination point agent type resource is configured with the value of the `FaultTolerance` attribute set to 0. At this value setting, the `CoordPoint` agent reports `FAULTED`, when the agent determines that at least one coordination point has keys (or registrations) missing.

The following is an example service group (`vxfen`) extracted from a `main.cf` file:

```
group vxfen (
  SystemList = { sysA = 0, sysB = 1 }
  AutoFailOver = 0
  Parallel = 1
  AutoStartList = { sysA, sysB }
)
  CoordPoint coordpoint (
    FaultTolerance=0
  )
  // resource dependency tree
  //
  //   group vxfen
  //   {
  //     CoordPoint coordpoint
  //   }
```

Debug log levels

The `CoordPoint` agent uses the following debug log levels:

`DBG_1`, `DBG_10`

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.

This agent is Intelligent Monitoring Framework (IMF)-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about IMF and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

The agent needs an exact match of PathName and Arguments for all the VCS-managed processes in the process table. To clarify, the processes that the ps command lists should have exact match of PathName and Arguments for all the processes in the VCS configuration file `/etc/VRTSvcs/conf/config/main.cf`.

Note that the AMF kernel driver does not monitor kernel processes. Even if you have enabled intelligent monitoring for Process agent, you must use the traditional poll-based monitoring to monitor kernel processes.

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 a binary executable for the specified process (program.vfd)
- The existence of the same binary on all nodes (program.vfd)

For more information about using the high availability fire drill see the *Veritas Cluster Server Administrator'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 a process in the background with optional arguments and priority in the specified user context.
Offline	Terminates the process with a <code>SIGTERM</code> . If the process does not terminate, a <code>SIGKILL</code> 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.
<code>imf_init</code>	Initializes the agent to interface with the asynchronous monitoring framework (AMF) kernel driver. This function runs when the agent starts up.
<code>imf_getnotification</code>	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel driver. The agent continuously waits for notification and takes action on the resource upon notification.
<code>imf_register</code>	Registers the resource entities, which the agent must monitor, with the AMF kernel driver. For example, the function registers the PID for online monitoring of a process. This function runs for each resource after the resource goes into steady state (online or offline).
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	<p>Indicates that the specified process is running in the specified user context.</p> <p>The agent only reports the process as online if the value configured for PathName attribute exactly matches the process listing from the ps output.</p>
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 5-8 Required attribute

Required attribute	Description
PathName	<p>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.</p> <p>This attribute must not exceed 256 characters.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/usr/sbin/proc1"</p>

Table 5-9 Optional attributes

Optional attribute	Description
Arguments	<p>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.</p> <p>Type and dimension: string-scalar</p>

Table 5-9 Optional attributes

Optional attribute	Description
PidFile	<p>The file that contains the process ID for the monitoring process. Specify the PidFile attribute for the monitoring process to use the Pid. Otherwise, to complete the monitoring process the agent uses the ps output.</p> <p>Note that when you use scripts, or other indirect mechanisms, to start processes, you must set the PidFile attribute if the ps output is different from the configured values for the PathName or Arguments attributes.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/var/lock/sendmail.pid"</p>
Priority	<p>Priority that the process runs. Priority values range between -20 (highest) to +19 (lowest).</p> <p>Type and dimension: string-scalar</p> <p>Default: 10</p>
UserName	<p>This attribute is the owner of the process. The process runs with this user ID.</p> <p>Type and dimension: string-scalar</p> <p>Default: root</p>

Resource type definition

```

type Process (
    static keylist SupportedActions = { "program.vfd", getcksum }
    static str ArgList[] = { PathName, Arguments, UserName,
        Priority, PidFile }
    str PathName
    str Arguments
    str UserName = root
    str Priority = 10
    str PidFile
)

```

Sample configurations

Configuration for Linux

In this example, the Process agent starts, stops, and monitors sendmail. This process is started with two arguments as determined in the Arguments attribute. The pid stored in the PidFile attribute is used to monitor the sendmail process.

```
Process sendmail (  
  PathName = "/usr/sbin/sendmail"  
  Arguments = "-bd -q30m"  
  PidFile = "/var/run/sendmail.pid"  
)
```

Debug log levels

The Process agent uses the following debug log levels:

DBG_1, DBG_2, DBG_3, DBG_4, DBG_5

ProcessOnOnly agent

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

VCS uses this agent internally to monitor 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 5-10 Required attributes

Required attribute	Description
PathName	<p>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 PathName attribute must not exceed 256 characters.</p> <p>The value configured for this attribute needs to match the process listing from the ps output for the agent to display as ONLINE.</p> <p>Type and dimension: string-scalar</p>

Table 5-11 Optional attributes

Optional attribute	Description
Arguments	<p>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.</p> <p>Type and dimension: string-scalar</p> <p>Example: "-bd -q30m"</p>
IgnoreArgs	<p>A flag that indicates whether monitor ignores the argument list.</p> <ul style="list-style-type: none"> ■ 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. <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>

Table 5-11 Optional attributes

Optional attribute	Description
PidFile	<p>The file that contains the process ID for the monitoring process. Specify the PidFile attribute for the monitoring process to use the Pid. Otherwise, to complete the monitoring process the agent uses the ps output.</p> <p>Note that when you use scripts, or other indirect mechanisms, to start processes, you must set the PidFile attribute when the ps output is different from the configured values for the PathName or Arguments attributes.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/var/lock/sendmail.pid"</p>
Priority	<p>Priority with which the process will run. Priority values range between -20 (highest) to +19 (lowest).</p> <p>Type and dimension: string-scalar</p> <p>Default: 10</p>
UserName	<p>Owner of the process. The process runs with the user ID.</p> <p>Type and dimension: string-scalar</p> <p>Default: root</p>

Resource type definition

```

type ProcessOnOnly (
    static str ArgList[] = { PathName, Arguments, UserName,
    Priority, PidFile, IgnoreArgs }
    static str Operations = OnOnly
    str PathName
    str Arguments
    str UserName = root
    str Priority = 10
    str PidFile
    boolean IgnoreArgs = 0
)

```

Sample configurations

```

group VxSS (

```

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

Phantom phantom_vxss (
)

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

Debug log levels

The ProcessOnOnly agent uses the following debug log levels:
DBG_1, DBG_4

Infrastructure and support agents

This chapter contains the following agents:

- [“About the infrastructure and support agents”](#) on page 199
- [“NotifierMngr agent”](#) on page 200
- [“Proxy agent”](#) on page 207
- [“Phantom agent”](#) on page 211
- [“RemoteGroup agent”](#) on page 213

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 Administrator'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.

Other applications with the name `notifier` can interfere with the NotifierMngr agent. If `notifier` is started outside VCS control, VCS can only monitor the notifier process if its started with the absolute path. For example, use:

```
# /opt/VRTSvcs/bin/notifier -s m=xyz &
```

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 <code>SIGABORT</code> . If the process does not exit within one second, VCS sends a <code>SIGKILL</code> .
Monitor	Monitors the notifier process.
Clean	Sends <code>SIGKILL</code> .

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

Table 6-2 Optional attributes

Optional attribute	Description
EngineListeningPort	<p>Change this attribute if the VCS engine is listening on a port other than its default port.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 14141</p>

Table 6-2 Optional attributes

Optional attribute	Description
MessagesQueue	<p>Size of the VCS engine's message queue. Minimum value is 30.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 30</p>
NotifierListeningPort	<p>Any valid, unused TCP/IP port number.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 14144</p>
NotifierSourceIP	<p>If this attribute is populated, all the notifications sent from the notifier (SMTP and SNMP) will be sent from the interface having this IP address.</p> <p>Note: Make sure that the SourceIP given in this attribute is present in the /etc/hosts file or is DNS-resolvable.</p> <p>Type and dimension: string-scalar</p> <p>Example: "10.209.77.111"</p>
SmtptFromPath	<p>Set to a valid email address, if you want the notifier to use a custom email address in the FROM: field.</p> <p>Type and dimension: string-scalar</p> <p>Example: "usera@example.com"</p>
SmtptRecipients	<p>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.</p> <p>Note: SmtptRecipients is a required attribute if you specify SmtptServer.</p> <p>Type and dimension: string-association</p> <p>Example:</p> <p style="padding-left: 40px;">"james@example.com" = SevereError, "admin@example.com" = Warning</p>

Table 6-2 Optional attributes

Optional attribute	Description
SmtprReturnPath	<p>Set to a valid email address, if you want the notifier to use a custom email address in the Return-Path: <> field.</p> <p>If the mail server specified in SmtprServer does not support SMTP VRFY command, then you need to set the SmtprVrfyOff to 1 in order for the SmtprReturnPath value to take effect.</p> <p>Type and dimension: string-scalar Example: "usera@example.com"</p>
SmtprServerTimeout	<p>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.</p> <p>Type and dimension: integer-scalar Default: 10</p>
SmtprServerVrfyOff	<p>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 SmtprServer attribute while sending emails.</p> <p>Type and dimension: boolean-scalar Default: 0</p>
SnmprCommunity	<p>Specifies the community ID for the SNMP manager.</p> <p>Type and dimension: string-scalar Default: public</p>

Table 6-2 Optional attributes

Optional attribute	Description
SnmpdTrapPort	<p>Port on the SNMP console machine where SNMP traps are sent.</p> <p>If you specify more than one SNMP console, all consoles use this value.</p> <p>Type and dimension: string-scalar</p> <p>Default: 162</p>

Resource type definition

```

type NotifierMngr (
  static int RestartLimit = 3
  static str ArgList[] = { EngineListeningPort, MessagesQueue,
  NotifierListeningPort, NotifierSourceIP, SnmpdTrapPort,
  SnmpCommunity, SnmpConsoles, SmtServer, SmtServerVrfyOff,
  SmtServerTimeout, SmtReturnPath, SmtFromPath, SmtRecipients
  }
  int EngineListeningPort = 14141
  int MessagesQueue = 30
  int NotifierListeningPort = 14144
  str NotifierSourceIP
  int SnmpdTrapPort = 162
  str SnmpCommunity = public
  str SnmpConsoles{}
  str SmtServer
  boolean SmtServerVrfyOff = 0
  int SmtServerTimeout = 10
  str SmtReturnPath
  str SmtFromPath
  str SmtRecipients{}
)

```

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 (`vcadmin@example.com`).

Configuration

```
system north

system south

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

    Phantom my_phantom (
    )

    NIC    NicGrp_eth0 (
        Device = eth0
    )

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

    Proxy nicproxy(
        TargetResName = "NicGrp_eth0"
    )
```

```
NotifierMngr ntfr (  
    SnmpConsoles = { snmpserv = Information }  
    SntpServer = "smtp.example.com"  
    SntpRecipients = { "vcsadmin@example.com" =  
        SevereError }  
)  
  
ntfr requires nicproxy  
  
// resource dependency tree  
//  
//     group Grp1  
//     {  
//     NotifierMngr ntfr  
//         {  
//             Proxy nicproxy  
//         }  
//     }  
// }
```

Debug log levels

The NotifierMngr agent uses the following debug log levels:
DBG_1

Proxy agent

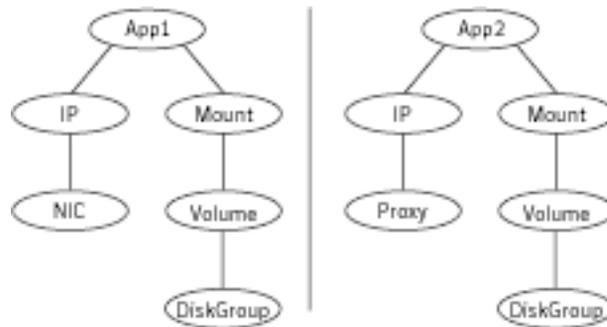
The Proxy agent mirrors the state of another resource on a local or remote system. It provides a means to specify and modify one resource and have its state reflected by its proxies. You can use the agent when you need to replicate the status of a resource.

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 that includes a Proxy resource



Agent functions

Monitor Determines status based on the target resource status.

Attributes

Table 6-3 Required attribute

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

Table 6-4 Optional attribute

Optional attribute	Description
TargetSysName	<p>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.</p> <p>Type and dimension: string-scalar</p> <p>Example: "sysa"</p>

Resource type definition

```

type Proxy (
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { TargetResName, TargetSysName,
        "TargetResName:Probed", "TargetResName:State" }
    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 3

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 @vcslx3 = { eth0 = "192.123.8.42", eth3 =
        "192.123.8.42" }
    Device @vcslx4 = { eth0 = "192.123.8.43", eth3 =
        "192.123.8.43" }
    NetMask = "255.255.248.0"
    NetworkHosts = { "192.123.10.129", "192.123.10.130" }
)

```

```
IPMultiNIC ip1 (  
  Address = "192.123.10.177"  
  MultiNICAResName = mnic  
  NetMask = "255.255.248.0"  
)
```

```
ip1 requires mnic
```

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

```
IPMultiNIC ip2 (  
  Address = "192.123.10.178"  
  NetMask = "255.255.255.0"  
  MultiNICAResName = mnic  
)
```

```
Proxy proxy (  
  TargetResName = mnic  
)
```

```
ip2 requires proxy
```

Debug log levels

The Proxy agent uses the following debug log levels:

DBG_1, DBG_2

Phantom agent

The Phantom agent enables VCS to determine the state of parallel service groups that do not include OnOff resources.

Do not use the Phantom resource in failover service groups.

Note: Do not attempt manual online or offline operations on the Phantom resource at the resource level. Do not use `hares` commands on the Phantom resource at the resource level. Unpredictable behavior results when you try a manual online or offline procedure or an `hares` command on a Phantom resource. You can perform commands on the service group that contains the Phantom resource.

Dependencies

No dependencies exist for the Phantom resource.

Figure 6-2 Sample service group that includes a Phantom resource



Agent functions

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

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 = 0, sysb = 1 }
  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:

- When it points to a global group
The RemoteGroup agent must then map the state of the global group in the local cluster.
- When it is configured inside a local parallel service group
The RemoteGroup resources on all cluster nodes monitor the same remote service group unless its attributes are localized.
- When it is configured inside a local failover service group

For more information on the functionality of this agent see the *Veritas Cluster Server Administrator'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 216 for more information.
Offline	Takes the remote service group offline. See the “ ControlMode ” on page 216 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 215.
Clean	If the RemoteGroup resource faults, the Clean function takes the remote service group offline. See the “ ControlMode ” on page 216 for more information.

State definitions

ONLINE	Indicates that the remote service group is in an ONLINE state. If the ReturnIntOffline attribute is not set to RemotePartial, then the remote service group is either in an ONLINE or PARTIAL state. See “ ReturnIntOffline ” on page 219.
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. The RemoteGroup resource returns intentional offline if the attribute ReturnIntOffline is set to an appropriate value. See “ ReturnIntOffline ” on page 219.
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 6-5 Required attributes

Required attribute	Description
IpAddress	<p>The IP address or DNS name of a node in the remote cluster. The IP address can be either physical or virtual.</p> <p>When configuring a virtual IP address of a remote cluster, do not configure the IP resource as a part of the remote service group.</p> <p>Type and dimension: string-scalar</p> <p>Examples: "www.example.com" or "11.183.12.214"</p>
Port	<p>This is a required attribute when the remote cluster listens on a port other than the default value of 14141.</p> <p>See "Port" on page 218.</p>
GroupName	<p>The name of the service group on the remote cluster that you want the RemoteGroup agent to monitor or manage.</p> <p>Type and dimension: string-scalar</p> <p>Example: "DBGp"</p>
VCSSysName	<p>You must set this attribute to either the VCS system name or the ANY value.</p> <ul style="list-style-type: none">■ ANY The RemoteGroup resource goes online if the remote service group is online on any node in the remote cluster.■ <i>VCSSysName</i> 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. <p>Type and dimension: string-scalar</p> <p>Example: "vcssys1" or "ANY"</p>

Table 6-5 Required attributes

Required attribute	Description
ControlMode	<p>Select only one of these values to determine the mode of operation of the RemoteGroup resource: MonitorOnly, OnlineOnly, or OnOff.</p> <ul style="list-style-type: none"> <p>■ 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.</p> <p>■ 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.</p> <p>■ 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.</p> <p>Type and dimension: string-scalar</p>

Table 6-5 Required attributes

Required attribute	Description
Username	<p>This is the login user name for the remote cluster.</p> <p>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.</p> <p>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.</p> <p>For a secure remote cluster:</p> <ul style="list-style-type: none"> ■ 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. <p>Type and dimension: string-scalar</p> <p>Example:</p> <ul style="list-style-type: none"> ■ For a cluster without the Symantec Product Authentication Service: "johnsmith" ■ For a secure remote cluster: "foobar@example.com"
Password	<p>This is the password that corresponds to the user that you specify in the Username attribute. You must encrypt the password with the <code>vcscrypt -agent</code> command.</p> <p>Note: Do not use the <code>vcscrypt</code> utility when entering passwords from a configuration wizard or the Cluster Manager (Java Console).</p> <p>Type and dimension: string-scalar</p>

Table 6-6 Optional attributes

Optional attribute	Description
DomainType	<p>For a secure remote cluster only, enter the domain type information for the specified user.</p> <p>For users who have the domain type unixpwd, you do not have to set this attribute.</p> <p>Type: string-scalar</p> <p>Example: "nis", "nisplus"</p>
BrokerIp	<p>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.</p> <p>Type: string-scalar</p> <p>Example: "128.11.295.51"</p>
Port	<p>The port where the remote engine listens for requests.</p> <p>This is an optional attribute, unless the remote cluster listens on a port other than the default value of 14141.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 14141</p>
OfflineWaitTime	<p>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.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p>

Table 6-6 Optional attributes

Optional attribute	Description
ReturnIntOffline	<p>Select one of the following values for RemoteGroup to return IntentionalOffline:</p> <ul style="list-style-type: none"> ■ RemotePartial—Indicates that the RemoteGroup resource returns an IntentionalOffline if the remote service group is in an ONLINE PARTIAL state. ■ RemoteOffline—Indicates that the RemoteGroup resource returns an IntentionalOffline if the remote service group is in an OFFLINE state. ■ RemoteFaulted—Indicates that the RemoteGroup resource returns an IntentionalOffline if the remote service group is OFFLINE FAULTED. <p>You can use these values in combinations with each other.</p> <p>You must set the IntentionalOffline attribute of the RemoteGroup resource type to 1 for this attribute to work properly. For more information about this attribute, see the <i>Veritas Cluster Server Administrator's Guide</i>.</p> <p>Type and dimension: string-vector Default: ""</p>
OfflineMonitoringNode	<p>Defines the cluster node that performs the offline monitoring of the remote service group.</p> <p>This is an internal attribute. Do not modify.</p>

Table 6-7 Type-level attributes

Type level attributes	Description
OnlineRetryLimit OnlineWaitLimit	<p>In case of remote service groups that take a longer time to Online, Symantec recommends that you modify the default OnlineWaitLimit and OnlineRetryLimit attributes.</p> <p>See the <i>Veritas Cluster Server Administrator's Guide</i> for more information about these attributes.</p>

Table 6-7 Type-level attributes

Type level attributes	Description
ToleranceLimit MonitorInterval	If you expect the RemoteGroup agent to tolerate sudden offlines of the remote service group, then modify the ToleranceLimit attribute. See the <i>Veritas Cluster Server Administrator's Guide</i> for more information about these attributes.
ExternalStateChange	If you want the local service group to go online or offline when the RemoteGroup resource goes online or offline outside VCS control, set the attribute ExternalStateChange appropriately. See the <i>Veritas Cluster Server Administrator's Guide</i> for more information about these attributes.

Resource type definition

```

type RemoteGroup (
    static int OnlineRetryLimit = 2
    static int ToleranceLimit = 1
    static boolean IntentionalOffline = 1
    static str ArgList[] = { IPAddress, Port, Username, Password,
        GroupName, VCSSysName, ControlMode, OfflineWaitTime,
        DomainType, BrokerIp, ReturnIntOffline }
    str IPAddress
    int Port = 14141
    str Username
    str Password
    str GroupName
    str VCSSysName
    str ControlMode
    int OfflineWaitTime
    str DomainType
    str BrokerIp
    str ReturnIntOffline[] = {}
    temp str OfflineMonitoringNode
)

```

Debug log levels

The RemoteGroup agent uses the following debug log levels:

DBG_1

Testing agents

This chapter contains the following agents:

- [“About the testing agents”](#) on page 221
- [“ElifNone agent”](#) on page 222
- [“FileNone agent”](#) on page 224
- [“FileOnOff agent”](#) on page 226
- [“FileOnOnly agent”](#) on page 228

About the testing agents

Use the testing agents to provide high availability for program support resources. These resources are useful for testing service groups.

ElifNone agent

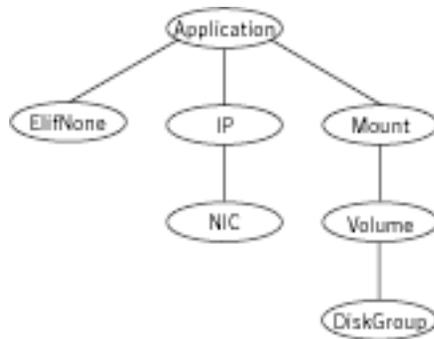
The ElifNone agent monitors a file. It checks for the file's absence.

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 that includes 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.
---------	--

State definitions

ONLINE	Indicates that the file specified in the PathName attribute does not exist.
FAULTED	Indicates that the file specified in the PathName attribute exists.
UNKNOWN	Indicates that the value of the PathName attribute does not contain a file name.

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"  
)
```

Debug log levels

The ElifNone agent uses the following debug log levels:

DBG_1, DBG_4, DBG_5

FileNone agent

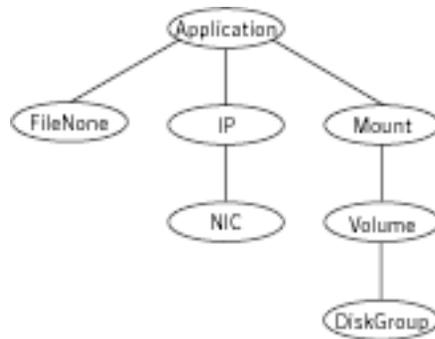
Monitors a file—checks for the file’s existence.

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 that includes 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.
---------	--

State definitions

ONLINE	Indicates that the file specified in the PathName attribute exists.
FAULTED	Indicates that the file specified in the PathName attribute does not exist.
UNKNOWN	Indicates that the value of the PathName attribute does not contain a file name.

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

Sample configuration

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

Debug log levels

The FileNone agent uses the following debug log levels:

DBG_1, DBG_4, DBG_5

FileOnOff agent

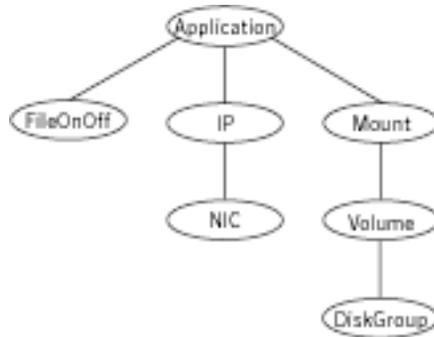
The FileOnOff agent creates, removes, and monitors files.

You can use this 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 that includes 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.

State definitions

ONLINE	Indicates that the file specified in the PathName attribute exists.
OFFLINE	Indicates that the file specified in the PathName attribute does not exist.

FAULTED	Indicates that the file specified in the PathName attribute has been removed out of VCS control.
UNKNOWN	Indicates that the value of the PathName attribute does not contain a file name.

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"  
)
```

Debug log levels

The FileOnOff agent uses the following debug log levels:
DBG_1, DBG_4, DBG_5

FileOnOnly agent

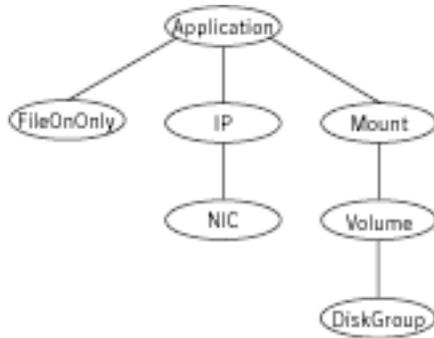
The FileOnOnly agent creates and monitors files.

You can use this 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 that includes 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.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	Indicates that the file specified in the PathName attribute exists.
OFFLINE	Indicates that the file specified in the PathName attribute does not exist and VCS has not attempted to bring the resource online.

FAULTED	Indicates that the file specified in the PathName attribute has been removed out of VCS control.
UNKNOWN	Indicates that the value of the PathName attribute does not contain a file name.

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"  
)
```

Debug log levels

The FileOnOnly agent uses the following debug log levels:

DBG_1, DBG_4, DBG_5

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 administrative IP address of the system.

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

Numerics

802.1Q trunking 80

A

about

- Network agents 79
- Samba agents 145

agent

modifying 18

agent functions

- Apache Web server agent 163
- Application agent 178
- CoordPoint agent 186
- DiskGroup agent 23
- DiskGroupSnap agent 33
- DiskReservation agent 45
- DNS agent 116
- ElifNone agent 222
- FileNone agent 224
- FileOnOff agent 226
- FileOnOnly agent 228
- IP agent 82
- IPMultiNIC agent 94
- LVMLogicalVolume agent 56
- LVMVolumeGroup agent 59
- Mount agent 65
- MultiNICA agent 104
- NetBIOS agent 157
- NFS agent 129
- NFSRestart agent 134
- NIC agent 88
- NotifierMngr agent 200
- Phantom agent 211
- Process agent 191
- ProcessOnOnly agent 195
- Proxy agent 207
- RemoteGroup agent 214
- SambaServer agent 148
- SambaShare agent 153
- Share agent 141
- Volume agent 50

VolumeSet agent 53

agents

- Apache Web server 162
- Application 176
- CoordPoint 186
- DiskGroup 22
- DiskGroupSnap 32
- DiskReservation 45
- DNS 115
- ElifNone 222
- FileNone 224
- FileOnOff 226
- FileOnOnly 228
- IP 81
- IPMultiNIC 94
- LVMLogicalVolume 56
- LVMVolumeGroup 59
- Mount 65
- MultiNICA 102
- NetBIOS 156
- NFS 128
- NFSRestart 133
- NIC 87
- NotifierMngr 200
- Phantom 211
- Process 190
- ProcessOnOnly 195
- Proxy 207
- RemoteGroup 213
- SambaServer 148
- SambaShare 153
- Share 140
- Volume 50
- VolumeSet 53

agents, typical functions 17

Apache Web server agent

- agent functions 163
- attributes 164
- description 162
- detecting application failure 170
- sample configuration 171
- state definitions 163

- Application agent
 - agent functions 178
 - attributes 180
 - description 176
 - high availability fire drill 176
 - resource type definition 183
 - sample configurations 184
 - state definitions 179
- association dimension 18
- attribute data types 18
- attributes
 - Application agent 180
 - CoordPoint agent 187
 - DiskGroup agent 25
 - DiskGroupSnap agent 34
 - DiskReservation agent 46
 - DNS agent 118
 - ElifNone agent 223
 - FileNone agent 225
 - FileOnOff agent 227
 - FileOnOnly agent 229
 - IP agent 83
 - IPMultiNIC agent 96
 - LVMLogicalVolume agent 57
 - LVMVolumeGroup agent 60
 - Mount agent 69
 - MultiNICA agent, 104
 - NFS agent 130
 - NFSRestart agent 136
 - NIC agent 89
 - NotifierMngr agent 201
 - Process agent 192
 - ProcessOnOnly 196
 - Proxy agent 208
 - RemoteGroup agent 215
 - SambaServer agent 150
 - Share agent 142
 - Volume agent 51
 - VolumeSet agent 54
- attributes, modifying 17, 18

B

- bonded network interfaces 88
- boolean data types 18
- bundled agents 17

C

- Cluster Manager (Java Console), modifying

- attributes 18
- CNAME record 123
- configuration files
 - main.cf 212
 - modifying 18
 - types.cf 17
- configuring, Samba agents 146
- CoordPoint agent
 - agent functions 186
 - attributes 187
 - description 186
 - resource type definition 188
 - sample configurations 189
 - state definitions 187

D

- data type
 - boolean 18
 - string 18
- data types
 - integer 18
- description, resources 17
- dimensions
 - keylist 18
 - scalar 18
 - vector 18
- DiskGroup agent
 - agent functions 23
 - attributes 25
 - description 22
 - high availability fire drill 29
 - resource type definition 29
 - sample configurations 30
 - state definitions 24
- DiskGroupSnap agent
 - agent functions 33
 - attributes 34
 - description 32
 - resource type definition 37
 - sample configurations 38
 - state definitions 33
- DiskReservation agent
 - agent functions 45
 - attributes 46
 - description 45
 - resource type definition 48
 - sample configurations 49
 - state definitions 45
- DNS agent 117

- agent functions 116
- attributes 118
- description 115
- resource type definition 122
- sample web server configuration 123

E

- ElifNone agent
 - agent functions 222
 - attributes 223
 - description 222
 - resource type definition 223
 - sample configuration 223
 - state definitions 222

F

- FileNone agent
 - agent functions 224
 - attribute 225
 - description 224
 - resource type definition 225
 - sample configurations 225
 - state definitions 224
- FileOnOff agent
 - agent functions 226
 - attribute 227
 - description 226
 - state definitions 226
- FileOnOnly agent
 - agent functions 228
 - attribute 229
 - description 228
 - resource type definition 229
 - sample configuration 229
 - state definitions 228

H

- high availability fire drill 29, 76, 81, 87, 123, 137, 176, 190

I

- integer data types 18
- IP agent
 - agent functions 82
 - attributes 83
 - description 81
 - high availability fire drill 81

- resource type definitions 85
- sample configurations 86
- state definitions 82

IPMultiNIC agent

- agent functions 94
- attributes 96
- description 94
- resource type definitions 98
- sample configuration 98
- state definitions 95

K

- keylist dimension 18

L

- LVMLogicalVolume agent
 - agent functions 56
 - attributes 57
 - description 56
 - resource type definition 57
 - sample configurations 57
 - state definitions 56
- LVMVolumeGroup agent
 - agent functions 59
 - attributes 60
 - description 59
 - resource type definition 61
 - sample configurations 62
 - state definitions 60

M

- main.cf 17, 212
- main.xml 17
- modifying
 - configuration files 18
- modifying agents 18
- monitor scenarios, DNS agent 123
- monitoring bonded NICs, Linux 91
- Mount agent
 - agent functions 65, 68
 - attributes 69
 - description 65
 - high availability fire drill 76, 123, 137
 - notes 76
 - resource type definition 75
 - sample configurations 78
- MultiNICA agent

- agent functions 104
- attributes 104
- description 102
- IP Conservation mode 103
- Performance mode 103
- resource type definitions 110
- resource type definitions, Linux 110
- sample configurations 110

MultiNICA agent, Linux

- IP Conservation mode 103
- Performance Mode 103

N

NetBIOS agent

- agent functions 157
- description 156
- resource type definition 158
- sample configurations 160
- state definitions 157

NFS agent

- agent functions 129
- attributes 130
- description 128
- resource type definition 131
- sample configurations 132
- state definitions 129

NFS lock recovery 131

NFSRestart agent

- agent functions 134
- attributes 136
- description 133
- resource type definition 137
- sample configuration 138
- state definitions 135

NIC agent

- agent functions 88
- attributes 89
- description 87
- high availability fire drill 87
- resource type definitions 90
- sample configurations 92
- state definitions 88

noautoimport flag 30

NotifierMngr agent

- agent functions 200
- attributes 201
- description 200
- resource type definition 204
- sample configurations 205

- state definitions 200

O

- online query 123

P

Phantom agent

- agent functions 211
- description 211
- resource type definition 211
- sample configurations 211

prerequisites

- NFS lock recovery 131
- Samba agents 145

Process agent

- agent functions 191
- attributes 192
- description 190
- high availability fire drill 190
- resource type definition 193
- sample configurations 194
- state definitions 192

ProcessOnOnly agent

- agent functions 195
- attributes 196
- description 195
- resource type definition 197
- sample configurations 197
- state definitions 195

Proxy agent

- agent functions 207
- attributes 208
- description 207
- resource type definition 209
- sample configurations 209

R

RemoteGroup agent

- agent functions 214
- attributes 215
- description 213
- resource type definition 220
- state definitions 214

resource type definition 52

- SambaShare agent 154

resource type definitions

- Application agent 183

- CoordPoint agent 188
- DiskGroup agent 29
- DiskGroupSnap agent 37
- DiskReservation agent 48
- DNS agent 122
- ElifNone agent 223
- FileNone agent 225
- FileOnOnly agent 229
- IP agent 85
- IPMultiNIC agent 98
- LVMLogicalVolume agent 57
- LVMVolumeGroup agent 61
- Mount agent 75
- MultiNICA agent 110
- MultiNICA agent, Linux 110
- NetBIOS agent 158
- NFS agent 131
- NFSRestart agent 137
- NIC agent 90
- NotifierMngr agent 204
- Phantom agent 211
- Process agent 193
- ProcessOnOnly agent 197
- Proxy agent 209
- RemoteGroup agent 220
- SambaServer agent 152
- Share agent 143
- Volume agent 52
- VolumeSet agent 54
- resource types 17
- resources
 - description of 17

S

- Samba agents 145
 - overview 145
 - prerequisites 145
- Samba agents configuring 146
- SambaServer agent
 - agent functions 148
 - attributes 150
 - description 148
 - resource type definition 152
 - sample configuration 152
 - state definitions 149
- SambaShare agent 153
 - agent functions 153
 - attributes 154
 - resource type definition 154

- sample configurations 155
 - state definitions 153
- sample configurations
 - Apache Web server agent 171
 - Application agent 184
 - CoordPoint agent 189
 - DiskGroup agent 30
 - DiskGroupSnap agent 38
 - DiskReservation agent 49
 - ElifNone agent 223
 - FileNone agent 225
 - FileOnOff agent 227
 - FileOnOnly agent 229
 - IP agent 86
 - IPMultiNIC 98
 - LVMLogicalVolume agent 57
 - LVMVolumeGroup agent 62
 - Mount agent 78
 - MultiNICA agent 110
 - NetBIOS agent 160
 - NFS agent 132
 - NFSRestart agent 138
 - NIC agent 92
 - NotifierMngr agent 205
 - Phantom agent 211
 - Process agent 194
 - ProcessOnOnly agent 197
 - Proxy agent 209
 - SambaServer agent 152
 - SambaShare agent 155
 - Share agent 144
 - Volume agent 52
- scalar dimension 18
- secure DNS update 124
- setting Mii and miimon 91
- Share agent
 - agent functions 141
 - attributes 142
 - description 140
 - resource type definitions 143
 - sample configurations 144
 - state definitions 141
- state definitions 54, 117
 - Apache Web server agent 163
 - Application agent 179
 - CoordPoint agent 187
 - DiskGroup agent 24
 - DiskGroupSnap agent 33
 - DiskReservation agent 45

- DNS agent 117
- ElifNone agent 222
- FileNone agent 224
- FileOnOff agent 226
- FileOnOnly agent 228
- IP agent 82
- IPMultiNIC agent 95
- LVMLogicalVolume agent 56
- LVMVolumeGroup agent 60
- Mount agent 68
- NetBIOS agent 157
- NFS agent 129
- NFSRestart agent 135
- NIC agent 88
- NotifierMngr agent 200
- Process agent 192
- ProcessOnOnly agent 195
- RemoteGroup agent 214
- SambaServer agent 149
- SambaShare agent 153
- Share agent 141
- Volume agent 50
- VolumeSet agent 54
- string data type 18

T

- trunking 80
- types.cf 17

V

- VCS, resource types 17
- vector dimension 18
- Volume agent
 - agent functions 50
 - attributes 51
 - description 50
 - sample configurations 52
 - state definitions 50
- VolumeSet agent 54
 - agent functions 53
 - attributes 54
 - description 53
 - resource type definition 54