

# Veritas™ Cluster Server Bundled Agents Reference Guide

Linux

6.0

# 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: 6.0

Document version: 6.0.3

## Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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 software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

[www.symantec.com/business/support/index.jsp](http://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](http://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/](http://www.symantec.com/business/support/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/business/support/](http://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](mailto:customercare_apac@symantec.com)

Europe, Middle-East, and Africa [semea@symantec.com](mailto:semea@symantec.com)

North America and Latin America [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

## Documentation

Product guides are available on the media 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.

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

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

[doc\\_feedback@symantec.com](mailto:doc_feedback@symantec.com)

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

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

## About Symantec Connect

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

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

# Contents

Technical Support .....	4
Chapter 1	Introducing Bundled agents ..... 17
	About Bundled agents ..... 17
	Resources and their attributes ..... 18
	Modifying agents and their resources ..... 18
	Attributes ..... 18
	IMF aware agents ..... 19
	Enabling debug log messages ..... 19
	VCS support for multi-pathing solutions ..... 20
Chapter 2	Storage agents ..... 22
	About the storage agents ..... 22
	DiskGroup agent ..... 22
	Dependencies for DiskGroup agent ..... 23
	Agent functions for DiskGroup agent ..... 23
	State definitions for DiskGroup agent ..... 25
	Attributes for DiskGroup agent ..... 25
	Resource type definition for DiskGroup agent ..... 29
	Notes for DiskGroup agent ..... 30
	Sample configurations for DiskGroup agent ..... 32
	Debug log levels for DiskGroup agent ..... 32
	DiskGroupSnap agent ..... 32
	Dependencies for DiskGroupSnap agent ..... 33
	Agent functions for DiskGroupSnap agent ..... 34
	State definitions for DiskGroupSnap agent ..... 34
	Attributes for DiskGroupSnap agent ..... 34
	Notes for DiskGroupSnap agent ..... 36
	Resource type definition for DiskGroupSnap agent ..... 40
	Sample configurations for DiskGroupSnap agent ..... 40
	Debug log levels for DiskGroupSnap agent ..... 47
	DiskReservation agent ..... 47
	Agent functions for DiskReservation agent ..... 47
	State definitions for DiskReservation agent ..... 47
	Attributes for DiskReservation agent ..... 48

Resource type definition for DiskReservation agent .....	49
Notes for DiskReservation agent .....	49
Sample configurations for DiskReservation agent .....	50
Debug log levels for DiskReservation agent .....	51
Volume agent .....	51
Dependencies for Volume agent .....	51
Agent functions for Volume agent .....	52
State definitions for Volume agent .....	52
Attributes for Volume agent .....	53
Resource type definition for Volume agent .....	53
Sample configuration for Volume agent .....	53
Debug log levels for Volume agent .....	54
VolumeSet agent .....	54
Dependencies for VolumeSet agent .....	54
Agent functions for VolumeSet agent .....	54
State definitions for VolumeSet agent .....	55
Attributes for VolumeSet agent .....	55
Resource type definition for VolumeSet agent .....	55
Sample configurations for VolumeSet agent .....	56
Agent notes for VolumeSet agent .....	56
Inaccessible volumes prevent the VolumeSet agent from coming online .....	56
Debug log levels for VolumeSet agent .....	56
LVMLogicalVolume agent .....	56
Dependencies for LVMLogicalVolume agent .....	57
Agent functions for LVMLogicalVolume agent .....	57
State definitions for LVMLogicalVolume agent .....	57
Attributes for LVMLogicalVolume agent .....	58
Resource type definition for LVMLogicalVolume agent .....	58
Sample configurations for LVMLogicalVolume agent .....	58
Debug log levels for LVMLogicalVolume agent .....	59
LVMVolumeGroup agent .....	60
Dependencies for LVMVolumeGroup agent .....	60
Agent functions for LVMVolumeGroup agent .....	60
State definitions for LVMVolumeGroup agent .....	61
Attributes for LVMVolumeGroup agent .....	61
Resource type definition for LVMVolumeGroup agent .....	62
LVMVolumeGroup agent notes .....	62
Enabling volume group activation protection for Veritas Dynamic Multi-Pathing .....	62
Notes for volume group activation protection .....	63
Sample configurations for LVMVolumeGroup agent .....	63
Debug log levels for LVMVolumeGroup agent .....	65



Mount agent .....	65
Dependencies for Mount agent .....	65
Agent functions for Mount agent .....	66
State definitions for Mount agent .....	68
Attributes for Mount agent .....	69
Resource type definition for Mount agent .....	73
Notes for Mount agent .....	74
Support for spaces in directory names .....	74
High availability fire drill .....	75
VxFS file system lock .....	75
IMF usage notes .....	76
Enabling Level two monitoring for the Mount agent .....	76
Sample configurations for Mount agent .....	77
Debug log levels for Mount agent .....	77

## Chapter 3

Network agents .....	78
About the network agents .....	78
Agent comparisons .....	78
IP agent .....	79
High availability fire drill for IP agent .....	80
Dependencies for IP agent .....	80
Agent functions for IP agent .....	80
State definitions for IP agent .....	81
Attributes for IP agent .....	81
Resource type definition for IP agent .....	84
Sample configurations for IP agent .....	85
Debug log levels for IP agent .....	85
NIC agent .....	85
High availability fire drill for NIC agent .....	86
Dependencies for NIC agent .....	86
Bonded network interfaces for NIC agent .....	86
Agent functions for NIC agent .....	87
State definitions for NIC agent .....	87
Attributes for NIC agent .....	88
Resource type definition for NIC agent .....	90
Monitoring bonded NICs for NIC agent .....	90
Setting Mii and miimon for NIC agent .....	90
Case 1 .....	90
Case 2 .....	91
Case 3 .....	91
Sample configurations for NIC agent .....	91
Debug log levels for NIC agent .....	92

IPMultiNIC agent .....	92
Dependencies for IPMultiNIC agent .....	92
Agent functions for IPMultiNIC agent .....	93
State definitions for IPMultiNIC agent .....	93
Attributes for IPMultiNIC agent .....	94
Resource type definition for IPMultiNIC agent .....	95
Sample configuration: IPMultiNIC and MultiNICA .....	96
Debug log levels .....	99
MultiNICA agent .....	99
Dependencies for MultiNICA agent .....	100
IP Conservation Mode (ICM) for MultiNICA agent .....	100
Performance Mode (PM) for MultiNICA agent .....	101
Agent function for MultiNICA agent .....	101
Attributes for MultiNICA agent .....	102
Resource type definition for MultiNICA agent .....	107
Sample configurations for MultiNICA agent .....	108
IPv6 configuration for MultiNICA agent .....	111
Mixed mode configuration—IPv4 and IPv6 for MultiNICA agent .....	112
Debug log levels for MultiNICA agent .....	113
DNS agent .....	113
Dependencies for DNS agent .....	113
Agent functions for DNS agent .....	114
State definitions for DNS agent .....	115
Attributes for DNS agent .....	116
Resource type definition for DNS agent .....	122
Agent notes for DNS agent .....	122
Sample configurations for DNS agent .....	127
Debug log levels for DNS agent .....	129
 Chapter 4	
File share agents .....	130
About the file service agents .....	130
NFS agent .....	130
Dependencies for NFS agent .....	131
Agent functions for NFS agent .....	131
State definitions for NFS agent .....	132
Attributes for NFS agent .....	132
Resource type definition for NFS agent .....	133
Notes for NFS agent .....	133
Sample configurations for NFS agent .....	134
Debug log levels for NFS agent .....	134
NFSRestart agent .....	135

Dependencies for NFSRestart agent .....	135
Agent functions for NFSRestart agent .....	136
State definitions for NFSRestart agent .....	137
Attributes for NFSRestart agent .....	138
Resource type definition for NFSRestart agent .....	138
Notes for NFSRestart agent .....	139
Sample configurations for NFSRestart agent .....	140
Debug log levels for NFSRestart agent .....	141
Share agent .....	141
Dependencies for Share agent .....	141
Agent functions for Share agent .....	142
State definitions for Share agent .....	142
Attributes for Share agent .....	143
Resource type definition for Share agent .....	144
Notes for Share agent .....	144
Sample configurations for Share agent .....	145
Debug log levels for Share agent .....	145
About the Samba agents .....	145
The Samba agents .....	146
Before using the Samba agents .....	146
Supported versions for the Samba agents .....	146
Notes for configuring the Samba agents .....	146
SambaServer agent .....	148
Dependencies for SambaServer agent .....	148
Agent functions for SambaServer agent .....	148
State definitions for SambaServer agent .....	149
Attributes for SambaServer agent .....	149
Resource type definitions for SambaServer agent .....	151
Sample configurations for SambaServer agent .....	151
Debug log levels for SambaServer agent .....	151
SambaShare agent .....	152
Dependencies for SambaShare agent .....	152
Agent functions for SambaShare agent .....	152
State definitions for SambaShare agent .....	152
Attributes for SambaShare agent .....	153
Resource type definition for SambaShare agent .....	153
Sample configuration for SambaShare agent .....	153
Debug log levels for SambaShare agent .....	154
NetBios agent .....	154
Dependencies for NetBios agent .....	154
Agent functions for NetBios agent .....	155
State definitions for NetBios agent .....	155
Attributes for NetBios agent .....	156

	Resource type definition for NetBios agent .....	157
	Sample configuration for NetBios agent .....	157
	Debug log levels for NetBios agent .....	158
Chapter 5	Service and application agents .....	159
	About the services and applications agents .....	159
	Apache HTTP server agent .....	159
	Dependencies .....	160
	Agent functions .....	160
	State definitions .....	161
	Attributes .....	162
	Resource type definition .....	165
	Apache HTTP server notes .....	166
	Sample configurations .....	169
	Application agent .....	173
	High availability fire drill for Application agent .....	174
	Dependencies for Application agent .....	174
	Agent functions for Application agent .....	175
	State definitions for Application agent .....	177
	Attributes for Application agent .....	178
	Resource type definition for Application agent .....	181
	Notes for Application agent .....	182
	Sample configurations for Application agent .....	183
	Debug log levels for Application agent .....	183
	CoordPoint agent .....	184
	Dependencies for CoordPoint agent .....	184
	Agent functions for CoordPoint agent .....	185
	State definitions for CoordPoint agent .....	185
	Attributes for CoordPoint agent .....	186
	Resource type definition for CoordPoint agent .....	186
	Notes for the CoordPoint agent .....	186
	Sample configuration for CoordPoint agent .....	187
	Debug log levels for CoordPoint agent .....	188
	KVMGuest agent .....	188
	Dependencies for KVMGuest agent .....	188
	Agent functions for KVMGuest agent .....	189
	State definitions for KVMGuest agent .....	190
	Attributes for KVMGuest agent .....	190
	Resource type definition for KVMGuest agent .....	192
	Notes for KVMGuest agent .....	192
	Sample configurations for KVMGuest agent .....	193
	Debug log levels for KVMGuest agent .....	195

Process agent .....	195
High availability fire drill for Process agent .....	196
Dependencies for Process agent .....	196
Agent functions for Process agent .....	196
State definitions for Process agent .....	197
Attributes for Process agent .....	198
Resource type definition for Process agent .....	199
Usage notes for Process agent .....	199
Sample configurations for Process agent .....	199
Debug log levels for Process agent .....	200
ProcessOnOnly agent .....	200
Dependencies .....	200
Agent functions .....	200
State definitions .....	201
Attributes .....	201
Resource type definition .....	202
ProcessOnOnly agent usage notes .....	203
Sample configurations .....	203
Debug log levels .....	203

Chapter 6	Infrastructure and support agents .....	204
	About the infrastructure and support agents .....	204
	NotifierMngr agent .....	204
	Dependency .....	205
	Agent functions .....	205
	State definitions .....	205
	Attributes .....	206
	Resource type definition .....	208
	Sample configuration .....	209
	Debug log levels .....	210
	Proxy agent .....	211
	Dependencies .....	211
	Agent functions .....	211
	Attributes .....	212
	Resource type definition .....	212
	Sample configurations .....	212
	Debug log levels .....	214
	Phantom agent .....	214
	Dependencies .....	214
	Agent functions .....	215
	Resource type definition .....	215
	Sample configurations .....	215

RemoteGroup agent .....	216
Dependency .....	216
Agent functions .....	217
State definitions .....	217
Attributes .....	218
Resource type definition .....	223
Debug log levels .....	223
Chapter 7      Testing agents .....	224
About the testing agents .....	224
ElifNone agent .....	224
Dependencies for ElifNone agent .....	224
Agent function for ElifNone agent .....	225
State definitions for ElifNone agent .....	225
Attributes for ElifNone agent .....	225
Resource type definition for ElifNone agent .....	226
Sample configuration for ElifNone agent .....	226
Debug log levels for ElifNone agent .....	226
FileNone agent .....	226
Dependencies for FileNone agent .....	226
Agent functions for FileNone agent .....	227
State definitions for FileNone agent .....	227
Attributes for FileNone agent .....	228
Resource type definition for FileNone agent .....	228
Sample configuration for FileNone agent .....	228
Debug log levels for FileNone agent .....	228
FileOnOff agent .....	228
Dependencies for FileOnOff agent .....	229
Agent functions for FileOnOff agent .....	229
State definitions for FileOnOff agent .....	229
Attribute for FileOnOff agent .....	230
Resource type definition for FileOnOff agent .....	230
Sample configuration for FileOnOff agent .....	230
Debug log levels for FileOnOff agent .....	230
FileOnOnly agent .....	230
Dependencies for FileOnOnly agent .....	230
Agent functions for FileOnOnly agent .....	231
State definitions for FileOnOnly agent .....	231
Attribute for FileOnOnly agent .....	232
Resource type definition for FileOnOnly agent .....	232
Sample configuration for FileOnOnly agent .....	232
Debug log levels for FileOnOnly agent .....	232

Chapter 8	Replication agents .....	233
	About the replication agents .....	233
	RVG agent .....	233
	Dependencies for RVG agent .....	234
	Agent functions for RVG agent .....	235
	State definitions for RVG agent .....	235
	Attributes for RVG agent .....	235
	Resource type definitions for RVG agent .....	236
	Sample configurations for RVG agent .....	236
	RVGPrimary agent .....	237
	Dependencies for RVGPrimary agent .....	237
	Agent functions for RVGPrimary agent .....	238
	State definitions for RVGPrimary agent .....	239
	Attributes for RVGPrimary agent .....	239
	Resource type definitions for RVGPrimary agent .....	242
	Sample configurations for RVGPrimary agent .....	243
	RVGSnapshot .....	243
	Dependencies for RVGSnapshot .....	244
	Agent functions for RVGSnapshot .....	244
	State definitions for RVGSnapshot .....	244
	Attributes for RVGSnapshot .....	245
	Resource type definitions for RVGSnapshot .....	245
	Sample configurations for RVGSnapshot .....	246
	RVGShared agent .....	246
	Dependencies for RVGShared agent .....	246
	Agent functions for RVGShared agent .....	247
	State definitions for RVGShared agent .....	247
	Attributes for RVGShared agent .....	248
	Resource type definitions for RVGShared agent .....	248
	Sample configurations for RVGShared agent .....	248
	RVGLogowner agent .....	248
	Dependencies for RVGLogowner agent .....	249
	Agent functions RVGLogowner agent .....	249
	State definitions RVGLogowner agent .....	250
	Attributes RVGLogowner agent .....	250
	Resource type definitions RVGLogowner agent .....	251
	RVGLogowner agent notes RVGLogowner agent .....	251
	Sample configurations RVGLogowner agent .....	252
	RVGSharedPri agent .....	252
	Dependencies for RVGSharedPri agent .....	253
	Agent functions for RVGSharedPri agent .....	253
	State definitions for RVGSharedPri agent .....	254

Attributes for RVGSharedPri agent .....	254
Resource type definitions for RVGSharedPri agent .....	255
Sample configurations for RVGSharedPri agent .....	255
Index .....	256



# Introducing Bundled agents

This chapter includes the following topics:

- [About Bundled agents](#)
- [Resources and their attributes](#)
- [Modifying agents and their resources](#)
- [Attributes](#)
- [IMF aware agents](#)
- [Enabling debug log messages](#)
- [VCS support for multi-pathing solutions](#)

## About Bundled agents

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	Enclose strings, which are a sequence of characters, in double quotes (""). Optionally enclose strings in quotes when they begin with a letter, and contains only letters, numbers, dashes (-), and underscores (_).  A string can contain double quotes, but the quotes must be immediately preceded by a backslash. In a string, represent a backslash with two backslashes (\\).
integer	Signed integer constants are a sequence of digits from 0 to 9. You can precede them with a dash. They are base 10. Integers cannot exceed the value of a 32-bit signed integer: 2147483647.
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{}

## IMF aware agents

- Mount agent
- Process agent
- Application agent

## 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 for Mount agent”](#) on page 77.

## VCS support for multi-pathing solutions

This section applies to the following agents:

- LVMLogicalVolume agent
- LVMVolumeGroup agent

VCS supports Symantec Dynamic Multi-Pathing (DMP) that is included as a part of the Storage Foundation and High Availability (SFHA) suite of products. Symantec does not support multi-pathing solutions that are not explicitly listed in the hardware compatibility list (HCL). You can find the HCL on the SORT web site, under the Documentation tab. However, Symantec supports third-party solutions, which are included as a part of the operating systems.

Symantec aims to thoroughly test and support third-party and native solutions, but it is not possible to test all third-party multi-pathing applications. This is because of complex support matrix and a number of potential product combinations. Hence, Symantec does not officially support multi-pathing solutions that are not explicitly listed in the HCL. Also, advanced functionality such as I/O fencing with SCSI3-PGR is only supported with arrays and multi-pathing solutions listed in the HCL and only with Symantec Storage Foundation.

If you are using a third-party multi-pathing solution, Symantec understands your need of keeping data paths redundant and does not insist that you uninstall or disable the solution. Symantec does not consider third-party multi-pathing solutions as invalid and continues to troubleshoot any support issues. However, for persisting support issues related to multi-pathing solutions, you need to contact the multi-pathing vendor.

# Storage agents

This chapter includes the following topics:

- [About the storage agents](#)
- [DiskGroup agent](#)
- [DiskGroupSnap agent](#)
- [DiskReservation agent](#)
- [Volume agent](#)
- [VolumeSet agent](#)
- [LVMLogicalVolume agent](#)
- [LVMVolumeGroup agent](#)
- [Mount agent](#)

## About the storage agents

Storage agents monitor shared storage and make shared storage highly available. Storage includes shared disks, disk groups, volumes, and mounts.

## 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, See [“Notes for DiskGroup agent”](#) on page 30.

## Dependencies for DiskGroup agent

The DiskGroup resource does not depend on any other resources.

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



## Agent functions for DiskGroup agent

Online	Imports the disk group using the <code>vxdg</code> command.
Offline	Deports the disk group using the <code>vxdg</code> command.
Monitor	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.  <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.</p> <p>Initiate the info agent function by setting the InfoInterval 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 DiskType and FreeSize of the DiskGroup resource is:</p> <pre># hares -value diskgroupres ResourceInfo</pre>
	<p>Output includes:</p> <pre>DiskType sliced FreeSize 35354136</pre>
	<p>The value specified is in kilo bytes.</p>
Action	<p>Different action agent functions follow:</p> <ul style="list-style-type: none"><li>■ <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.</li><li>■ <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.</li><li>■ <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.</li><li>■ <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.</li><li>■ <code>volinuse</code> Checks if open volumes are in use or file systems on volumes that are mounted outside of VCS configuration.</li></ul> <p>See <a href="#">“High availability fire drill”</a> on page 30.</p>



## State definitions for DiskGroup agent

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 UseFence is not set to "SCSI3" but the Reservation attribute value is "SCSI3".

## Attributes for DiskGroup agent

**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	<p>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 service group containing the resource offline.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p><b>Note:</b> If the MonitorReservation attribute is set to 0, the value of the clusterwide attribute UseFence is set to SCSI3, and the disk group is imported without SCSI reservation, then the monitor agent function takes the service group containing the resource offline.</p>

**Table 2-2** Optional attributes (*continued*)

Optional attributes	Description
PanicSystemOnDGLoss	

Table 2-2 Optional attributes (*continued*)

Optional attributes	Description
	<p>Determines whether to panic the node if the disk group becomes disabled or monitor program times out. A loss of storage connectivity can cause the disk group to become disabled. VxVM commands not responding properly can cause monitor program to timeout.</p> <p><b>Note:</b> System administrators may want to set a high value for FaultOnMonitorTimeout to increase system tolerance.</p> <p>This attribute accepts following values 0, 1, 2,3</p> <ul style="list-style-type: none"> <li>■ 0 : Do not halt the system</li> <li>■ 1 : halt the system if either disk group goes into disabled state or the disk group resource faults due to monitor timeout</li> <li>■ 2: halt the system if disk group goes into disabled state</li> <li>■ 3: halt the system if disk group resource faults due to monitor timeout</li> </ul> <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"> <li>■ If the cluster has I/O fencing enabled, the DiskGroup resource is marked FAULTED. This state results in the agent attempting to take the service group offline. As part of bringing the DiskGroup resource offline, the agent attempts to deport the disabled disk group. Even if disabled disk group fails to deport, the DiskGroup resource enters a FAULTED state. This state enables the failover of the service group that contains the resource. To fail back the DiskGroup resource, manually deport the disk group after restoring storage connectivity.</li> <li>■ If the cluster does not use I/O fencing, a message is logged and the resource is reported ONLINE. The resource is reported ONLINE so that it does not fail over, which ensures data integrity.</li> </ul> <p><b>Note:</b> The PanicSystemOnDGLoss attribute does not depend on the MonitorReservation attribute.</p> <p><b>Note:</b> If PanicSystemOnDGLoss is set to non-zero value, the system panic is initiated using <code>poweroff</code></p>

Table 2-2 Optional attributes (continued)

Optional attributes	Description
	<p><code>-nf</code> command. This command halts the system. An administrator needs to bring up the system.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p>
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 <code>autostartvolumes</code> 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</p> <p>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</p> <p>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"><li>■ 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.</li><li>■ The agent attempts to forcefully unmount the file systems to close the volumes.</li></ul> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p>

**Table 2-2** Optional attributes (*continued*)

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"> <li>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.</li> <li>SCSI3—The disk group is imported with SCSI-3 reservation if the value of the cluster-level UseFence attribute is SCSI3.</li> <li>NONE—The disk group is imported without SCSI-3 reservation.</li> </ul> <p>Type and dimension: string-scalar Default: ClusterDefault Example: "SCSI3"</p>

**Table 2-3** Internal attributes

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 for DiskGroup agent

The resource definition for this agent on Linux follows:

```
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,  
ConfidenceLevel }  
str DiskGroup  
boolean StartVolumes = 1  
boolean StopVolumes = 1  
boolean MonitorReservation = 0  
int PanicSystemOnDGLoss = 0  
temp str tempUseFence = INVALID  
str DiskGroupType = private  
int UmountVolumes = 0  
str Reservation = ClusterDefault  
)
```

## Notes for DiskGroup agent

The DiskGroup agent has the following notes:

- [High availability fire drill](#)
- [Using volume sets](#)
- [Setting the noautoimport flag for a disk group](#)

### 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
```

## For VxVM version 4.0

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

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

To deport the disk group, enter:

```
# vxdg deport disk_group
```

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

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

## Sample configurations for DiskGroup agent

### DiskGroup resource configuration

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

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

## Debug log levels for DiskGroup agent

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.

---

**Note:** The DiskGroupSnap agent requires the Global Cluster Option (GCO) license enabled on all systems in the cluster.

---

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

You must define the DiskGroupSnap agent in a separate FireDrill service group which is similar to the Application service group. The FireDrill service group might contain resources similar to the Application service group, for example Mount, Application, and so on.

The FireDrill service group must also contain a resource of type DiskGroupSnap such that the Mount resource depends on the DiskGroupSnap resource. The main DiskGroup must contain multiple sites registered in it with the value of the "siteconsistent" attribute set to on.

When the DiskGroupSnap agent goes online, the agent detaches one of the sites from the main DiskGroup and imports the detached site on the fire drill host as an independent DiskGroup with a different name. The volumes on the DiskGroup are also imported and mounted with same names on the fire drill host.




The DiskGroupSnap agent provides Gold and Bronze configurations for the fire drill, which can be specified using the agent's FDType attribute. The configuration decides the site to be detached from the DiskGroup for fire drill testing. The Gold configuration is the default option in which the agent selects a site from the DiskGroup that is neither the local VxVM site nor the site on which the DiskGroup is online. With the Gold configuration, you can also specify an alternate site to detach through the agent's FDSiteName attribute. With the Bronze configuration, the agent uses the local VxVM site name as the site to detach from the DiskGroup.

For important information about this agent, See [“Notes for DiskGroupSnap agent”](#) on page 36.

## Dependencies for DiskGroupSnap agent

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 for DiskGroupSnap agent

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. After the completion of Online function, the agent creates a lock file in the lock directory (/var/VRTSvcs/lock) to indicate that the resource is online.
Offline	This re-attaches the site that the value of the FDSiteName attribute specifies back to the application's disk group. After the completion of Offline function the agent removes the lock file from the lock directory (/var/VRTSvcs/lock) to indicate that the resource is Offline.
Monitor	Monitors the DiskGroupSnap resource by checking the existence of the Lock file in /var/VRTSvcs/lock directory..
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 for DiskGroupSnap agent

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 for DiskGroupSnap agent

**Table 2-4** Required attributes

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

**Table 2-4** Required attributes (*continued*)

Required attribute	Description
FDType	<p>Specifies the configuration to be used for the fire drill. The possible values for this attribute are:</p> <ul style="list-style-type: none"> <li>■ Bronze</li> <li>■ Gold (default)</li> </ul> <p>The Bronze configuration uses the local host's VxVM site name as the site to be detached from the DiskGroup. This action leaves the DiskGroup vulnerable to site disaster since a copy of the production volume might not be available when the fire drill is in progress.</p> <p>In the Gold configuration there are at least three copies of the parent volume available on different sites, hence, even after detaching one site the volume is not vulnerable to site disaster while the fire drill is in progress.</p>

**Table 2-5** Optional attributes

Optional attribute	Description
FDSiteName	<p>The unique VxVM site name tag for the fire drill disks. The value of this attribute is used in conjunction with the FDType attribute and it must be set to one of the sites registered in the main DiskGroup.</p> <ul style="list-style-type: none"> <li>■ When FDType is set to the Bronze configuration, the value of FDSiteName should either be empty or the name of the local host VxVM site for the fire drill host.</li> <li>■ When FDType is set to the Gold configuration, FDSiteName identifies a site in the DiskGroup to detach as a part of the fire drill. If FDSiteName is left blank, the agent will choose a site to detach based on the DiskGroup configuration. The agent chooses a site name from the DiskGroup which is neither the production server's site name nor the fire drill host's site name.</li> </ul> <p><a href="#">Table 2-6</a> shows the possible values of the attributes FDType and FDSiteName and the decision taken by the agent.</p>

Consider a configuration where the Production DiskGroup contains three sites: A, B, and C, and the Application service group is online on a node with local VxVM site ID is A. Fire drill is being done on another node Application service group is online on a node where local VxVM site ID is B.

Table 2-6 Example FDType configurations

FDType	Bronze			Gold/Empty			
	FDSitename	Empty	B	C	Empty	B	C
Result	Use B as the site to detach and proceed	Detach site B from DiskGroup	Error	Check if there is another site other than A and B and select it. Else, it is an error	Error	Detach site C from the DiskGroup	

## Notes for DiskGroupSnap agent

The DiskGroupSnap agent has the following notes:

- See [“Fire drill configuration after upgrading VCS”](#) on page 36.
- See [“Configuring the SystemZones attribute for the fire drill service group”](#) on page 36.
- See [“Configuring the FireDrill service group”](#) on page 37.
- See [“Adding the ReuseMntPt attribute to the ArgList attribute for the Mount agent type”](#) on page 37.
- See [“Configuration considerations”](#) on page 38.
- See [“Agent limitations”](#) on page 39.

### Fire drill configuration after upgrading VCS

After upgrading VCS from any earlier version to 6.0, delete all resources of type DiskGroupSnap and recreate them again using the new definitions of the attributes. Failure to perform this step might result in an unexpected behavior of the agent.

### 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 fire drill 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 Netlsnr to 1:

```
# haconf -makerw
# hatype -modify Oracle FireDrill 1
# hatype -modify Netlsnr 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 fire drill 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 fire drill'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 fire drill 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`.

See [Figure 2-3](#) on page 39. shows the changes to resource values for the fire drill service group; note that the Volume resource is not included.

- Before commencing the fire drill, make sure that all the sites registered in the application DiskGroup are in ACTIVE state.

**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 fire drill service group offline before you shut down VCS on any node. If you fail to take the fire drill 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.
- Do not switch the Application service group when fire drill is in progress.

## Resource type definition for DiskGroupSnap agent

The resource type definition for this agent follows:

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

## Sample configurations for DiskGroupSnap agent

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 for DiskGroupSnap agent

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
)

FDtype = "Gold"
  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/newdgl_fd/newvol1"
  FSType = vxfs
  FsckOpt = "-y"
  ReuseMntPt = 1
)

Mount mntfdres2 (
  MountPoint = "/dgsfs2"
  BlockDevice = "/dev/vx/dsk/newdgl_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 = newdgl
)

Mount mntres1 (
    MountPoint = "/dgsfs1"
    BlockDevice = "/dev/vx/dsk/newdgl/newvol1"
    FSType = vxfs
    FsckOpt = "-y"
    ReuseMntPt = 1
)

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

Process procrs1 (
    PathName = "/usr/bin/ksh"
    Arguments = "/scrib.sh /dgsfs1"
)

Process procrs2 (
    PathName = "/usr/bin/ksh"
    Arguments = "/scrib.sh /dgsfs2"
)

mntres1 requires dgres
mntres2 requires dgres
```

```
procrs1 requires mntres1
procrs2 requires mntres2
```

## Sample main.cf of DiskGroupSnap with Oracle resource

The following Oracle configuration has been simplified for presentation within this guide.

```
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
    FDType = "Bronze"
)

IP fd_oraip (
    Device = eth0
    Address = "10.198.95.191"
    NetMask = "255.255.255.0"
)

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

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

NIC fd_oranic (
```

```
Device = eth0
NetworkHosts = { "10.198.95.1" }
)

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 = eth0
    Address = "10.198.95.192"
    Netmask = "255.255.252.0"
)

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 = eth0
)

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 for DiskGroupSnap agent

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, See [“Notes for DiskReservation agent”](#) on page 49.

## Agent functions for DiskReservation agent

The agent functions for this agent follow:

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 for DiskReservation agent

The state definitions for this agent follow:

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 for DiskReservation agent

**Table 2-7** 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: <code>"/dev/sdd"</code></p>

**Table 2-8** 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 2-8** Optional attributes (*continued*)

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 Corporation recommends a value is between 3 and 8.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 3</p>

## Resource type definition for DiskReservation agent

The resource definition for this agent on Linux follows:

```

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

```

## Notes for DiskReservation agent

The DiskReservation agent has the following notes:

- [The DiskReservation agent does not reserve disks that have multiple paths](#)

- [Configuring the MonitorTimeout attribute for more than three disks](#)

## 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 for DiskReservation agent

The sample configurations for this agent follow:

### Configuration 1 for DiskReservation agent

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/sdc"
    FSType = vxfs
    MountOpt = rw
)

mount requires diskres1

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

## Debug log levels for DiskReservation agent

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 for Volume agent

Volume resources depend on DiskGroup resources.

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

## Agent functions for Volume agent

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 for Volume agent

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 for Volume agent

**Table 2-9** Required attributes

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

**Table 2-10** Internal attribute

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

## Resource type definition for Volume agent

The resource type definition for this agent follows:

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

## Sample configuration for Volume agent

The sample configuration for the Volume agent follows:

```

Volume sharedg_vol3 (
    Volume = vol3
    DiskGroup = sharedg
)

```

## Debug log levels for Volume agent

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 for VolumeSet agent

VolumeSet resources depend on DiskGroup resources.

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



## Agent functions for VolumeSet agent

Online	Uses the vxrecover command to start the volume set.
Offline	Uses the vxvset 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 for VolumeSet agent

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 for VolumeSet agent

**Table 2-11** Required attributes

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

## Resource type definition for VolumeSet agent

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

## Sample configurations for VolumeSet agent

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 for VolumeSet agent

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 for VolumeSet agent

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.

See [“VCS support for multi-pathing solutions”](#) on page 20.

## Dependencies for LVMLogicalVolume agent

LVMLogicalVolume resources depend on LVMVolumeGroup resources.

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



## Agent functions for LVMLogicalVolume agent

Online	Starts the volume using the lvchange command.
Offline	Stops the volume using the lvchange 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 for LVMLogicalVolume agent

The state definitions for this agent follow:

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 for LVMLogicalVolume agent

**Table 2-12** Required attributes

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

## Resource type definition for LVMLogicalVolume agent

The resource definition for this agent on Linux follows:

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

## Sample configurations for LVMLogicalVolume agent

The sample configuration for the LVMLogicalVolume agent follows:

### Configuration 1 for LVMLogicalVolume agent

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

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

```
LVMVolumeGroup lvg1 (  
    VolumeGroup = testvg_1  
)  
DiskReservation dr1 (  
    Disks = { "/dev/sdc", "/dev/sdd" }  
)  
lvoll requires lvg1  
lvg1 requires dr1
```

## Configuration 2 for LVMLogicalVolume agent

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 for LVMLogicalVolume agent

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.


See [“VCS support for multi-pathing solutions”](#) on page 20.

### Dependencies for LVMVolumeGroup agent

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 that includes a LVMVolumeGroup resource



### Agent functions for LVMVolumeGroup agent

The agent functions for this agent follow:

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	Different action agent functions follow: <ul style="list-style-type: none"><li>■ volinuse 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 for LVMVolumeGroup agent

ONLINE	Indicates that the volume group is imported.
OFFLINE	Linux: 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 for LVMVolumeGroup agent

**Table 2-13** Required attributes

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

**Table 2-14** Optional attributes

Optional attribute	Description
StartVolumes	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.  Type and dimension: boolean-scalar Default: 0

**Table 2-14** Optional attributes (*continued*)

Optional attribute	Description
EnableLVMTagging	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.  Type and dimension: boolean-scalar  Default: 0

## Resource type definition for LVMVolumeGroup agent

The resource definition for this agent on Linux follows:

```
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 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 for LVMVolumeGroup agent

The sample configurations for this agent follow:

### 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" }  
)  
  
lvgl 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 lvgl (  
    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 for LVMVolumeGroup agent

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, and NFS.


---

This agent also supports high availability fire drills.

For important information about this agent, See [“Notes for Mount agent”](#) on page 74.

## Dependencies for Mount agent

The Mount resource does not depend on any other resources.

**Figure 2-9** Sample service group that includes a Mount resource

## Agent functions for Mount agent

Online	<p>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.</p> <p>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.</p>
Offline	Unmounts the mounted file system.
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). This action entry point registers mountpoint, blockdevice, and fstype for mount agent.
Clean	Unmounts the mounted file system forcefully.

## Info

The Mount agent info function executes the command:

```
# df -h mount_point
```

The output displays Mount resource information:

```
Size Used Avail Use%
```

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

```
# haconf -makerw
```

```
# hatype -modify Mount InfoInterval 60
```

The command to retrieve information about the Mount resource is:

```
# hares -value mountres ResourceInfo
```

Output includes:

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

## Action

- **chgmtlock**  
Resets the VxFS file system lock to a VCS-defined lock.
- **mountpoint.vfd**  
Checks if the specified mount point exists on the offline node. If it fails and you request that VCS fixes it, it creates the mount point directory using `mkdir` command.
- **mounted.vfd**  
Checks if the mount point is already mounted on the offline node. If it fails, you need to unmount all the file systems from the specified mount point directory.
- **vxslic.vfd**  
Checks for valid Veritas File System (VxFS) licenses. If it fails, you need to update the license for VxFS.
- **mountentry.vfd**  
Checks that the mount point is not listed in auto file system tables. For example,  
`/etc/fstab`  
If this action fails, you need to remove the mount point from auto file system tables.

## attr\_changed

Unlocks the mounts when you change the value of the VxFSMountLock attribute from 1 to 0.

## State definitions for Mount agent

The state definitions for this agent follow:

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 file system 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 file system 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 file system has unexpectedly unmounted.</p>
UNKNOWN	<p>Indicates that a problem exists either with the configuration or the ability to determine the status of the resource.</p>

## Attributes for Mount agent

**Table 2-15** 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>■ LVM2 example           <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 /dev/mapper/ora_vg-ora_vol for the BlockDevice attribute.</p> </li> <li>■ NFS example           <p>If the file system type is NFS, then specify the BlockDevice as:</p> <pre>server:/path/to/share galaxy.symantecexample.com:/usr/share1</pre> </li> <li>■ "/dev/vx/dsk/myvcs_dg/myvol"</li> <li>■ IPv6           <pre>"[fe80::1:2:3]:/dirname/anotherdir"</pre> </li> </ul>
FscckOpt	<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>. For example: <code>hares -modify MntRes FscckOpt %-y</code></p>

**Table 2-15** Required attributes (*continued*)

Required attribute	Description
FSType	Type of file system. Supports vxfs, bind, ext2, ext3,ext4, xfs, nfs, or reiserfs. Type and dimension: string-scalar
MountPoint	Directory for mount point. Type and dimension: string-scalar Example: "/mnt/apache1"

**Table 2-16** 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. If the value of this attribute is 0, and Storage Checkpoints are mounted, then failover does not occur. Type and dimension: boolean-scalar Default: 1
MountOpt	Options for the <code>mount</code> command. Refer to the <code>mount</code> manual page for more information. Do not specify <code>-o</code> in the MountOpt field. The agent uses this option only when bringing a Mount resource online. If the value of this attribute is <code>nfs</code> , the agent always mounts the file system with the "soft,intr" option. Type and dimension: string-scalar Example: "rw"

Table 2-16 Optional attributes (*continued*)

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>
SnapUmount	<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 2-16 Optional attributes (*continued*)

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</p> <p>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</p> <p>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</p> <p>Example: "usr1"</p>
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>



**Table 2-16** Optional attributes (*continued*)

Optional attribute	Description
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 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 Default: 0</p>
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 Default: 0</p>

## Resource type definition for Mount agent

The resource definition for this agent on Linux follows:

```
Type Definition
```

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

## Notes for Mount agent

The Mount agent has the following notes:

- [Support for spaces in directory names](#)
- [High availability fire drill](#)
- [VxFS file system lock](#)
- [IMF usage notes](#)
- [Enabling Level two monitoring for the Mount agent](#)

## Support for spaces in directory names

The Mount agent supports directory names with spaces. The space can be leading, trailing, or in the middle of the name. If the directory name has a trailing space,

provide an extra "/" at the end of the corresponding attribute of a Mount resource. The attributes that currently support spaces in directory names are MountPoint and BlockDevice. Note that the agent does not support spaces created using the TAB key.

## 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 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 and RHEL 6:
  - 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 Level two monitoring for the Mount agent

Level two monitoring can be enabled for the Mount agent only if FSType is set to “nfs”.

To enable Level two monitoring, run the following commands:

- # haconf -makerw
- # hares -override resource\_name LevelTwoMonitorFreq
- # hares -modify resource\_name LevelTwoMonitorFreq 1

- # haconf -dump -makero

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

## Sample configurations for Mount agent

### Basic SCSI configuration for Mount agent

Configuration for Linux follows:

```
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 for Mount agent

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

### NFS mount example for Mount agent

```
Mount mnt1 (  
MountPoint = "/mnt/vctest"  
BlockDevice = "sys1:/home/export"  
FSType = nfs  
)
```

## Debug log levels for Mount agent

The Mount agent uses the following debug log levels:

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

# Network agents

This chapter includes the following topics:

- [About the network agents](#)
- [IP agent](#)
- [NIC agent](#)
- [IPMultiNIC agent](#)
- [MultiNICA agent](#)
- [DNS agent](#)

## About the network agents

Use network agents to provide high availability for networking resources.

---

**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

Agent comparisons may be made as described in the following sections.

The network agents support IPv4 and IPv6 addresses.

### 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 Mode (ICM), 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 (ICM) or as soon as it comes up (PM)
- Use the original base IP address when failing over (ICM) or Require a pre-assigned base IP address for each NIC (PM)
- Have only one active NIC at a time in case of ICM and more than one active NIC at a time in case of PM

### 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 NIC and IP agents, VCS supports Linux bonded interfaces.

## High availability fire drill for IP agent

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 resource

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

## Dependencies for IP agent

IP resources depend on NIC resources.

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



## Agent functions for IP agent

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 ifconfig 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 ip command. For IPv6 addresses, the ip command is used.
Action	<ul style="list-style-type: none"> <li>■ route.vfd Checks for the existence of a route to the IP from the specified NIC.</li> <li>■ device.vfd 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 specified in the Address attribute.

## State definitions for IP agent

The state definitions for this agent follow:

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 for IP agent

For Linux:

**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.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>■ IPv4: "192.203.47.61"</li> <li>■ IPv6: "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: <code># sysctl -p /etc/sysctl.conf</code></p> <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>.</p> <p>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>
One of the following attribute: <ul style="list-style-type: none"> <li>■ PrefixLen</li> <li>■ NetMask</li> </ul>	See <a href="#">Table 3-2</a> on page 83.

**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). Configure this attribute 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 data-bbox="561 986 1206 1038">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>For complete list of <code>ifconfig</code> options refer to <i>ifconfig manpage</i>.</p> <p>Type and dimension: string-scalar</p> <p>Example: "broadcast 172.20.9.255"</p>

**Table 3-2** Optional attributes (*continued*)

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 -4 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 for IP agent

The resource definition for this agent on Linux follows:

```
type IP (
    static keylist RegList = { NetMask }
    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 for IP agent

The sample configurations for this agent follow:

### IPv4 Configuration

Configuration 1 for Linux follows:

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

### IPv6 Configuration

Configuration using a specified NetMask for Linux follows:

```
IP IP_2001_10 (  
Device = eth0  
Address = "2001::10"  
PrefixLen = 64  
)
```

## Debug log levels for IP agent

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 for NIC agent


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 for NIC agent

The NIC resource does not depend on any other resources.

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



## Bonded network interfaces for NIC agent

The NIC agent now supports using bonded network interfaces.

See [“Monitoring bonded NICs for NIC agent”](#) on page 90.

## Agent functions for NIC agent

### Monitor

- If the NIC maintains its connection status, the agent uses MII to determine the status of the resource.  
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 ONLINE.  
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 ONLINE. If the count remains unchanged, the agent sends a ping to the broadcast address of the device to generate traffic on the network.  
The agent counts the number of packets that the NIC receives before and after the broadcast. If the count increases, the resource is marked ONLINE. If the count remains the same or decreases over a period of five broadcast cycles, the resource faults.

### Action

- device.vfd  
Checks for the existence of the interface configured in the Device attribute.

---

**Note:** The NIC agent supports ethtool and MII-based device status monitoring. If the Mii attribute for a NIC resource is set to 1, the agent first checks for the device status using ethtool and if it fails, it checks for the MII status for the device. The agent does not do any ping and packet count-based monitoring if it can successfully decide the status of the NIC device based on ethtool and MII tests.

---

## State definitions for NIC agent

The state definitions for this agent follow:

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

## Attributes for NIC agent

**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 ethtool and 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 (*continued*)

Optional attribute	Description
NetworkHosts	<p>List of hosts on the network that are pinged 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> <p>For IPv4:</p> <pre># ping -I device hostip</pre> <p>For IPv6:</p> <pre># ping6 -I device 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> <p>IPv4:</p> <pre>{ "166.93.2.1", "166.99.1.2" }</pre> <p>IPv6:</p> <pre>{ "2001::1" , "166.93.2.1" }</pre>
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 for NIC agent

The resource definition for this agent on Linux follows:

```
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 for NIC agent

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 NIC agent

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 for NIC agent

### Configuration for using Mii for NIC agent

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 for NIC agent

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

### IPv6 configuration for NIC agent

The following is a basic configuration for IPv6 with IP and NIC resources.

```
group nic_group (  
    SystemList = { sysA = 0, sysB = 1 }  
    Parallel = 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" }
)
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 for NIC agent

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 the virtual IP to a backup interface. If multiple service groups have IPMultiNIC resources associated with the same MultiNICA resource, only one group must have the MultiNICA resource. The other groups have Proxy resources pointing to it. You can use this agent for IP addresses on multiple-adaptor systems.

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

## Dependencies for IPMultiNIC agent

IPMultiNIC resources depend on MultiNICA resources.

**Figure 3-3** Sample service group that includes an IPMultiNIC resource

## Agent functions for IPMultiNIC agent

Online	Configures a virtual IP address on the active interface of the MultiNICA resource. Also sends out a gratuitous ARP.
Offline	Removes the virtual IP address from the active interface of the MultiNICA resource.
Monitor	Checks if the virtual IP address is configured on one interface of the MultiNICA resource.
Clean	Removes the virtual IP address from active interface of the MultiNICA resource.

## State definitions for IPMultiNIC agent

The state definitions for this agent follow:

ONLINE	Indicates that the specified IP address is assigned to one of the interfaces specified in the corresponding MultiNICA resource.
OFFLINE	Indicates that the specified IP address is not assigned to any interface of the MultiNICA resource.
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 for IPMultiNIC agent

**Table 3-5** Required attributes

Required attribute	Description
Address	The virtual IP address that is assigned to the active NIC. Type and dimension: string-scalar Examples: <ul style="list-style-type: none"><li>■ IPv4: "10.128.10.14"</li><li>■ IPv6: "2001:DB8::"</li></ul>
MultiNICAResName	Name of the associated MultiNICA resource that determines the active NIC. Type and dimension: string-scalar Example: "mnic"
One of the two attributes: <ul style="list-style-type: none"><li>■ NetMask</li><li>■ PrefixLen</li></ul>	See <a href="#">Table 3-6</a> on page 94.

**Table 3-6** Optional attributes

Optional attribute	Description
Options	The <code>ifconfig</code> command options for the virtual IP address. Do not configure any options in this attribute that you can specify in other attributes. For example, for the netmask use the <code>NetMask</code> attribute. Type and dimension: string-scalar Example: "mtu 2000"

**Table 3-6** Optional attributes (*continued*)

Optional attribute	Description
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 "ip addr add <i>ipv4addr/prefixlen</i> <i>IPOptions</i> <i>label</i> <i>label</i> dev <i>device</i>"</li> <li>■ IPv6 "ip addr add <i>ipv6addr/prefixlen</i> <i>IPOptions</i> <i>label</i> <i>label</i> dev <i>device</i>"</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 ip command instead of ifconfig.</p> <p>Example:</p> <ul style="list-style-type: none"> <li>■ "broadcast 172.20.9.255"</li> <li>■ "scope link"</li> </ul>
NetMask	<p>For the IPv4 protocol, specify the value of NetMask attribute in decimal (base 10).</p> <p>Configure this attribute 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>Specifies the prefix for the IPv6 address represented as the CIDR value. 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>Note: 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>

## Resource type definition for IPMultiNIC agent

The resource definition for this agent on Linux follows:

```
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 for IPMultiNIC agent

```
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
// {
// MultiNIC mnic
// }
// }
// }
```

## IPv6 configuration for IPMultiNIC agent

```
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
)
MultiNIC mnic (
Device @sysA = { eth0 = "2001::10", eth1 = "2001::11" }
Device @sysB = { eth0 = "2001::12", eth2 = "2001::13" }
PrefixLen=96
NetworkHosts = { "2001::1", "2001::500" }
)
ip1 requires mnic
// resource dependency tree
//
// group grp1
// {
```

```
//      IPMultiNIC ip1
//      {
//      MultiNICA mnic
//      }
// }
```

## Mixed mode configuration—IPv4 and IPv6 for IPMultiNIC agent

Mixed mode configuration for IPv4 and IPv6 follows:

```
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 =
"192.123.11.128" }
Device @sysB = { eth0 = "192.123.10.129", eth2 =
"192.123.11.130" }
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\_2, 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-adaptor 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 \(ICM\) for MultiNICA agent](#)
- [Performance Mode \(PM\) for MultiNICA agent](#)

With sufficient IP addresses, use PM.


The IPMultiNIC and MultiNICA agents support Linux bonds.

The IPMultiNIC and MultiNICA agents support IPv4 and IPv6.

## Dependencies for MultiNICA agent

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) for MultiNICA agent

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

### Configuration for MultiNICA agent

When a MultiNICA resource is to be configured in ICM, the configured 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 for MultiNICA agent

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 tries to find the next working NIC in the order of priority.

## Performance Mode (PM) for MultiNICA agent

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 for MultiNICA agent

When the MultiNICA resource is to be configured in PM, 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 before starting VCS and handing over the management of the NICs to the agent.

## Operation for MultiNICA agent

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 for MultiNICA agent

The agent functions for this agent follow:

Monitor	Uses ethtool and Media 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 for MultiNICA agent

While configuring the MultiNICA resource, in addition to the 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 each system you must localize the attribute with separate base IP addresses.</p> <p>For IP Conservation Mode (ICM): 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           <ul style="list-style-type: none"> <li>Device@vcsonelnx1={ eth1 = "10.212.100.178", eth2 = "10.212.100.179" }</li> <li>Device@vcsonelnx2 = { eth2 = "10.212.100.180", eth3 = "10.212.100.181" }</li> </ul> </li> <li>■ IPv6           <ul style="list-style-type: none"> <li>Device@vcsonelnx1={ eth1 = "1234::5678", eth2 = "1234::5679" }</li> <li>Device@vcsonelnx2={ eth3 = "1234::5680", eth4 = "1234::5681" }</li> </ul> </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>
<p>One of the two attributes:</p> <ul style="list-style-type: none"> <li>■ NetMask</li> <li>■ PrefixLen</li> </ul>	

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>
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:</p> <pre>"ip addr add ipv4addr/prefixlen IPv4AddrOptions dev device"</pre> <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>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>



Table 3-8 Optional attributes (*continued*)

Optional attribute	Description
IPv4RouteOptions	<p>The extra options that are passed to the <code>ip route add</code> command for IPv4 addresses. The agent uses this attribute in tandem with the IPv4AddrOptions attribute. The <code>ip route add</code> command generally resembles:</p> <pre>"ip route add IPv4RouteOptions dev device"</pre> <p><b>Note:</b> If Options attribute is configured, and none of IPv4AddrOptions and LinkOptions is configured, 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>
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. The <code>ip addr add</code> command resembles:</p> <pre>"ip addr add ipv6addr/ prefixlen IPv6AddrOptions dev device"</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 IPv6RouteOptions device dev"</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>

Table 3-8 Optional attributes (*continued*)

Optional attribute	Description
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><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>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>
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>You must configure either this Options attribute or both the IPv4AddrOptions and IPv4RouteOptions attributes.</p> <p>This attribute does not support IPv6.</p> <p><b>Note:</b> If you configure any of the IPv4AddrOptions or LinkOptions attribute, the agent ignores Options attribute and uses <code>ip</code> command instead of <code>ifconfig</code>.</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 NetworkHosts 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>

Table 3-8 Optional attributes (*continued*)

Optional attribute	Description
RouteOptions	<p>Assignment of a base IP address to a device followed by a route add command. The command takes the options specified by this attribute. RouteOptions are applicable only when configuring the local host as the default gateway. No routes are added if this string is set to NULL.</p> <p>The RouteOptions attribute works in tandem with the Options attribute. You must configure the Options 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>
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 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>
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 if the IP address is an IPv4 address.</p> <p>Type and dimension: string-scalar</p> <p>Example: "255.255.252.0"</p>
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 ethtool and Mii to monitor the status of the NIC.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 1</p>

## Resource type definition for MultiNICA agent

```
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 for MultiNICA agent

The sample configurations for the following agent are:

### 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 =
"192.123.11.128" }
    Device @sysB = { eth0 = "192.123.10.129", eth2 =
"192.123.11.130" }
    NetMask = "255.255.248.0"
    NetworkHosts = { "192.123.10.1", "192.123.10.2" }
)

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 = { "192.123.10.1", "192.123.10.2" }
    Failback = 0

)

ip1 requires mnic

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

## IPv6 configuration for MultiNICA agent

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" }
    Failback = 0
)
```

```
        PrefixLen = 64
    )

    ipmnic_res requires mnic_res
```

## Mixed mode configuration—IPv4 and IPv6 for MultiNICA agent

Mixed mode configuration for IPv4 and IPv6 follows:

```
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
  Failback = 0
  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" }
  Failback=0
  PrefixLen=96
  NetworkHosts = { "2001::1", "192.123.10.1" }
```



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

## Debug log levels for MultiNICA agent

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)
- Alias to hostname or 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 if the Resource Records need to be dynamically added and deleted from the DNS servers during failover. The agent updates the name server with the new resource record mappings while failing over and allows the clients to connect to the failed over instance of the application.

For important information about this agent, refer to [Agent notes for DNS agent](#)

## Dependencies for DNS agent

No dependencies exist for the DNS resource.

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

## Agent functions for DNS agent

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> <p><b>Note:</b> The DNS agent does not send any update for a resource record if it is already present on the name server.</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> <p>The monitor entry point also sends periodic refresh requests to DNS server if the RefreshInterval attribute is set.</p>
Clean	<p>Removes the Online lock file, if it exists. If attribute OffDeIRR is true, clean removes all records that the ResRecord keys define.</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 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 This action entry point checks if dig and nsupdate binaries exist and are executable.</li><li>■ master.vfd This action entry point checks if stealth masters are able to reply to SOA query for the configured domain.</li></ul>

## State definitions for DNS agent

The state definitions for this agent follow:

ONLINE	<p>Online lock file exists and at least one name server can return all configured resource records.</p>
OFFLINE	<p>At least one of the following is true:</p> <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	<p>Indicates that the DNS resource is not configured correctly. 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.</p>

## Attributes for DNS agent

**Table 3-9** Required attributes

Required attribute	Description
Domain	<p>A string representing the DNS zone that the agent administers.</p> <p>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: "demo.example.com"</li><li>■ IPv4 reverse mapping: "2.168.192.in-addr.arpa"</li></ul>

**Table 3-9** Required attributes (*continued*)

Required attribute	Description
ResRecord	

**Table 3-9** Required attributes (*continued*)

Required attribute	Description
	<p>ResRecord is an association of DNS resource record values. Each ResRecord attribute consists of two values: <i>DNS record key = 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. The A type ResRecord configuration should be as follows: <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 }</li> <li>or</li> <li>■ ResRecord = { www = "mydesktop.marketing.example.com." }</li> </ul> <p>Where the Domain attribute is "marketing.example.com"</p> <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 types are as</p>

**Table 3-9** Required attributes (*continued*)

Required attribute	Description
	<p>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 an IPv4 address (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 the agent logs a warning message to the engine log file. This 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 3-10** Optional attributes

Optional attribute	Description
TTL	<p>This attribute (a non-zero integer) represents the Time To Live (TTL) 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 TTL 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 3-10** Optional attributes (*continued*)

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. This attribute should be configured only when the DNS server configured is a Unix based DNS server.</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 functions 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 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>
OffDelRR	<p>Use the OffDelRR attribute to direct the offline and clean agent functions to remove all records that the ResRecord key defines. You must set the value of this attribute to 1 (true) to have the agent remove all the records.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>



**Table 3-10** Optional attributes (*continued*)

Optional attribute	Description
UseGSSAPI	<p>Use the UseGSSAPI attribute if the DNS server that you have configured is a Windows DNS server and only if it accepts secure dynamic updates.</p> <p><b>Note:</b> Do not set this attribute if the Windows DNS server accepts non-secure updates.</p> <p>If this attribute is set to 1, the agent uses the -g option with the nsupdate command.</p> <p>See <a href="#">"Agent notes for DNS agent"</a> on page 122. for more information on requirements to use the DNS agent with the secure Windows DNS server.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>
RefreshInterval	<p>This attribute represents the time interval in seconds after which the DNS agent attempts to refresh the resource records (RRs) on the DNS servers. The default value of zero indicates that the DNS agent does not attempt to refresh the records on the DNS servers. The DNS agent writes the warning message to the logs if it is not able to refresh the DNS records.</p> <p><b>Note:</b> The refresh request is sent in the next monitor cycle after the RefreshInterval period is reached.</p> <p>If the DNS agent is unable to refresh the DNS records, and the records are removed as a result of a scavenging operation or by the DNS administrator, the DNS resource will fault.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p> <p>Example: 3600</p>

**Table 3-10** Optional attributes (*continued*)

Optional attribute	Description
CleanRRKeys	<p>Use this attribute to direct the online agent function to clean up all the existing DNS records for the configured keys before adding new records. The default value (0) disables this behavior.</p> <p><b>Note:</b> If multiple DNS resources are configured with the same key value in their ResRecord attribute, then do not set this attribute value to 1.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>

## Resource type definition for DNS agent

The resource definition for this agent on Linux follows:

```

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

```

## Agent notes for DNS agent

The DNS agent has the following notes:

- [About using the VCS DNS agent on UNIX with a secure Windows DNS server](#)
- [High availability fire drill for DNS agent](#)

- [Monitor scenarios for DNS agent](#)
- [Sample Web server configuration for DNS agent](#)
- [Secure DNS update for BIND 9 for DNS agent](#)
- [Setting up secure updates using TSIG keys for BIND 9 for DNS agent](#)

## About using the VCS DNS agent on UNIX with a secure Windows DNS server

This section describes the requirements for using the DNS agent with a secure Windows DNS server. Note that there are no special requirements for sending non-secure updates to a Windows DNS server.

### Software requirement for DNS agent

For the secure updates on Windows DNS server to work, the VCS DNS agent on UNIX requires BIND version 9.7.2-P3 or later installed on all cluster nodes.

### Configuration requirement for DNS agent

The VCS DNS agent on UNIX requires setting up Kerberos authentication with the Windows DNS server and configuring the domain and DNS server information in `/etc/resolv.conf` at the client node.

To set up the Kerberos authentication from the UNIX host to the Windows DNS server, configure the Kerberos configuration file (`/etc/krb5.conf` or `/etc/krb/krb5.conf`) to use the Windows DNS server as Key Distribution Centre (KDC).

A sample Kerberos configuration file with domain `privdns.sym` and DNS server `master.privdns.sym` is as follows:

```
[libdefaults]
default_realm = PRIVDNS.SYM
dns_lookup_realm = true
dns_lookup_kdc = true
default_tkt_enctypes = des-cbc-md5
default_tgs_enctypes = des-cbc-md5
ticket_lifetime = 24h
renew_lifetime = 7d
forwardable = true
allow_weak_crypto = true
[realms]
PRIVDNS.SYM = {
kdc = master.privdns.sym:88
kpasswd_server = master.privdns.sym:464
admin_server = master.privdns.sym
```

```
}  
[domain_realm]  
.privdns.sym = PRIVDNS.SYM  
privdns.sym = PRIVDNS.SYM
```

---

**Note:** Symantec does not support KDC and Domain Controller/DNS located on different servers.

---

Authenticate all the nodes on the cluster (on which the DNS agent is configured to run) with the Active directory. Use kinit on your user account and use klist to verify that you have a ticket to the configured realm principal. Refer to the man page of kinit for more information on obtaining Kerberos ticket granting tickets from KDC.

---

**Note:** The DNS agent requires a node to be authenticated with Kerberos all the time. Renew the obtained tickets periodically if your authentication method requires you to do so.

---

A sample run of kinit and klist for the above configuration with user vcsdns will look as follows:

```
# kinit vcsdns  
Password for vcsdns@PRIVDNS.SYM:  
# klist  
Ticket cache: FILE:/tmp/krb5cc_0  
Default principal: vcsdns@PRIVDNS.SYM  
Valid starting Expires Service principal  
12/14/09 16:17:37 12/15/09 02:19:09 krbtgt/PRIVDNS.SYM@PRIVDNS.SYM  
renew until 12/21/09 16:17:37
```

If the environment variable KRB5CCNAME is set to some non-default location (default is /tmp), then VCS will not inherit it by default and will look for the Kerberos tickets in default location /tmp.

To resolve this issue, un-set the environment variable KRB5CCNAME and run the kinit command again. This will update the Kerberos tickets in default location (/tmp). Else, for a customized location (for example, /cache/krb\_ticket) for Kerberos tickets, add an entry in opt/VRTSvcs/bin/vcsenv file on each cluster node before VCS starts:

```
KRB5CCNAME="FILE:/cache/krb_ticket"  
export KRB5CCNAME
```

Update `/etc/resolv.conf` on your client node to add information for the Windows DNS server and the configured domain.

## High availability fire drill for DNS agent

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 for DNS agent

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 3-11](#) summarizes the monitor scenarios for the Online lock files.

**Table 3-11** 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 for DNS agent

Take the example of a Web server. A browser requests the URL `http://www.example.com` that maps to the canonical name `server1.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 server one to server two (`server2.example.com`), the domain name servers need a new canonical name

mapping for `www.example.com`. After the failover, the DNS resource updates this mapping of `www.example.com` to point to canonical name `server2.example.com`

---

**Note:** In this configuration, the Domain attribute should be configured with value "example.com"

---

## Secure DNS update for BIND 9 for DNS agent

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 BIND 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 for DNS agent

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

To use secure updates using TSIG keys, perform the following steps at the DNS server:

- 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 HOST example.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 are 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 for DNS agent

This section contains sample configurations for this agent.

### Basic IPv6 configuration for DNS agent

This sample configuration provides basic configuration for IPv6 support. In the following configuration, `nic_value` represents the base NIC value for the platform

For example: `eth0`

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

```
DNS ipv6group_dns_res (
    Critical = 0
    Domain = "example.com"
    TSIGKeyFile = "/var/tsig/Kipv6.vcsd.net.+157+18435.private"
    StealthMasters = { "2001:db8:c18:2:69c4:3251:bac1:6cbe" }
    ResRecord = {
        vcssysCv6 = "2001:db8:c18:2:214:4fff:fe96:8833",
        sysC = vcssysCv6 }
    )

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 for DNS agent

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 = { www = server1 }
    ResRecord @sysB = { www = server2 }
    OffDelRR = 1
    )
```



## IPv4 A sample configuration for DNS agent

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 for DNS agent

The DNS agent uses the following debug log levels:

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

# File share agents

This chapter includes the following topics:

- [About the file service agents](#)
- [NFS agent](#)
- [NFSRestart agent](#)
- [Share agent](#)
- [About the Samba agents](#)
- [SambaServer agent](#)
- [SambaShare agent](#)
- [NetBios agent](#)

## 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.

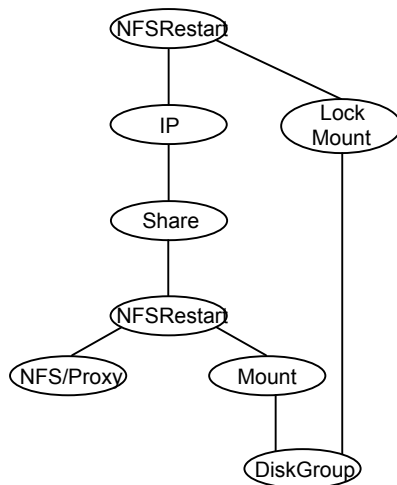
For important information about this agent,

See [“Notes for NFS agent”](#) on page 133.

## Dependencies for NFS agent

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 for NFS agent

Online	Starts NFS services on the system. For NFSv3, the services also include NSM (Network Status Monitor) and NLM (Network Lock Manager) service. If NFSv4 is enabled, it also starts NFSv4 name mapping daemon
Monitor	Monitors the NFS services running on the system. For NFSv3, it checks for version 2 and version 3 of NFS service along with other NFS services. For NFSv4, it also checks the availability of NFSv4 name mapping daemon.
Clean	Stops and restarts NFS services for all kernels. It also stops and restarts NSM and NLM services running on the system.

## State definitions for NFS agent

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 for NFS agent

**Table 4-1** Optional attributes for Linux

Optional attributes	Description
GracePeriod	Required when the value of the NFSRestart attribute is 1. GracePeriod specifies the amount of time that lock recovery is allowed by the NFS server after its reboot.  Type and dimension: integer-scalar Default: 90
LockFileTimeout	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.  Type and dimension: integer-scalar Default: 180
Nproc	Specifies the number of concurrent NFS requests that the server can handle.  Type and dimension: integer-scalar Default: 8 Example: 16

**Table 4-1** Optional attributes for Linux (*continued*)

Optional attributes	Description
NFSSecurity	<p>Specifies whether to start the NFS security daemon <code>rpc.svcgssd</code> or not. You must configure the type of security that NFS supports, for example: Kerberos.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>
NFSv4Support	<p>Specifies whether to start the NFSv4 daemon <code>rpc.idmapd</code> or not and whether to monitor <code>nfsd</code> version 4.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>

## Resource type definition for NFS agent

```

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
)

```

## Notes for NFS agent

The NFS agent has the following notes:

- [Prerequisites for NFS lock recovery](#)
- [Using NFSv4](#)

### 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. Only one of the Share resources that depends on the NFS resource needs a value of fsid=0 in its Options attribute. 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 just 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 for NFS agent

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 for NFS agent

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.

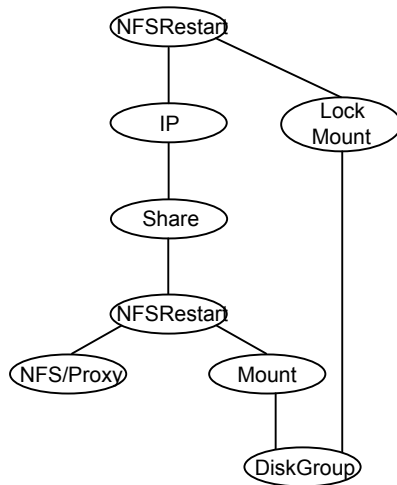
If you have configured the NFSRestart agent for lock recovery, the NFSRestart agent starts the smsyncd daemon. The daemon copies the NFS locks from the local directory `/var/lib/nfs` to shared storage. The agent's online function copies the locks from shared storage to local directory `/var/lib/nfs`.

For important information about this agent, refer to [Notes for NFSRestart agent](#)

## Dependencies for NFSRestart agent

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. Both the NFSRestart resources provide combined protection from potential corruption of NFS locks and potential NFS ACK storms. 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 for NFSRestart agent

The agent functions for this agent follow:

### 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>nfscnf.vfd 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 Verifies that the NFS lock directory (which is specified by the LocksPathName attribute of NFSRestart) is on shared storage.</li></ul>

## State definitions for NFSRestart agent

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 for NFSRestart agent

**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.  Type and dimension: string-scalar 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.  Type and dimension: string-scalar Example: "/share1x"
NFSLockFailover	NFS Lock recovery is done for all the Share resources that are configured in the group of this resource.  Type and dimension: boolean-scalar Default: 0
Lower	Defines the position of NFSRestart resource in the service group. The NFSRestart resource below the Share resource needs a value of 1.  The NFSRestart resource on the top of the resource dependency tree has a Lower attribute value of 0.  Type and dimension: integer-scalar Default: 0

## Resource type definition for NFSRestart agent

```
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  
)
```

## Notes for NFSRestart agent

The NFSRestart agent has the following notes:

- [About high availability fire drill](#)
- [Providing a fully qualified host name](#)

### 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 for NFSRestart agent

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 for NFSRestart agent

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.

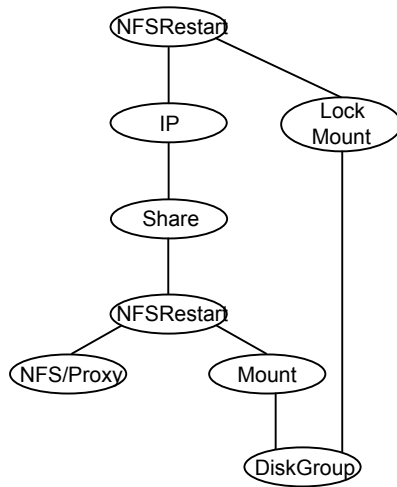
For important information on this agent, refer to:

[Notes for Share agent](#)

## Dependencies for Share agent

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 for Share agent

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	<p>direxists.vfd</p> <p>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.</p>

## State definitions for Share agent

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 specified directory is unshared outside the control of VCS.

## Attributes for Share agent

**Table 4-4** Required attributes

Required attribute	Description
PathName	Pathname of the file system to be shared. Type and dimension: string-scalar 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'.  Client or host where the directory specified by PathName is exported. The client can be a wild card (*), a fully qualified domain name (FQDN) including the host name, or an IP address. For more information on specifying IP address, refer to <a href="#">About entering an IP address in the Client attribute</a> .  Type and dimension: string-scalar 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:  "rw, no_root_squash"  For more information about the <code>exportfs</code> command and its options, refer to the <code>exportfs</code> manpage.  Type and dimension: string-vector Default = "ro, async, wdelay, root_squash"

**Table 4-5** Optional attributes (*continued*)

Optional attribute	Description
OtherClients	<p>The Client attribute can be assigned one FQDN host name or IP address, whereas multiple FQDN host names, or IP addresses can be assigned to the <b>OtherClients</b> field. For more information on specifying IP address, refer to <a href="#">About entering an IP address in the Client attribute</a>.</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 for Share agent

```

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
)

```

## Notes for Share agent

The following section contains notes on the Share agent.

- [Support for spaces in directory names](#)
- [High availability fire drill](#)
- [About entering an IP address in the Client attribute](#)

### Support for spaces in directory names

The Share agent supports directory names with spaces. The space can be leading, trailing, or in the middle of the name. If the directory name has a trailing space, provide an extra "/" at the end of the PathName attribute of a Share resource. Note that the agent does not support spaces created using the TAB key.



## 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*.

## About entering an IP address in the Client attribute

You can enter an IP address as a part of the Client or OtherClients attribute. If you specify an IP address, use the form `a.b.c.d/32`. Using this form of IP address prevents the monitor entry point from reporting a Share resource as offline even if the sharetab entry is present. This occurs on some SLES systems where the `exportfs` command tries to resolve the client IP address to the host name and rewrite the sharetab entry with the host name. If you specify the IP address as `a.b.c.d/32`, `exportfs` does not recognize this as an IP address and does not attempt to resolve the address to a host name.

## Sample configurations for Share agent

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 for Share agent

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 for the Samba agents

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

```
# smbd -V
```

## 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 because 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. Only one resource of this type is permitted. You can use the agent to make a `smbd` daemon highly available.

The `smbd` daemon provides Samba share services. The agent verifies that Samba is running by reading the pid of `smbd` 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 for SambaServer agent

No dependencies exist for the SambaServer resource.

**Figure 4-4** Sample service group that includes a SambaServer resource



## Agent functions for SambaServer agent

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 pid file. Does indepth 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 forcefully if required.

## State definitions for SambaServer agent

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 for SambaServer agent

**Table 4-6** Required attributes

Required attribute	Description
ConfFile	Complete path of the configuration file that Samba uses. Type and dimension: string-scalar Example: "/etc/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. Type and dimension: string-scalar Example: "/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. Type and dimension: integer-scalar Default: 5

Table 4-7 Optional attributes (*continued*)

Optional attribute	Description
Ports	<p>Ports where Samba accepts connections.</p> <p>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.</p> <p>Type and dimension: integer-vector</p> <p>Default: 139, 445</p>
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>
PidFile	<p>The absolute path to the Samba daemon pid file. This file contains the process ID of the monitored <code>smbd</code> process.</p> <p>Configure this attribute if you are using a non-standard configuration file name or path. If this attribute 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: <code>"/var/run/smbd.pid"</code></p>
SambaTopDir	<p>Parent path of Samba daemon and binaries.</p> <p>Type and dimension: string-scalar</p> <p>Example: <code>"/usr"</code></p> <p>SambaServer agent uses SambaTopDir attribute value in an open entry point to determine the complete path of samba executables. If this attribute is configured after the resource is enabled, please disable and enable the resource again to bring this into effect as follows:</p> <pre># hares -modify &lt;res&gt; Enabled 0 # hares -modify &lt;res&gt; Enabled 1</pre>

**Table 4-7** Optional attributes (*continued*)

Optional attribute	Description
SocketAddress	<p>The IP address where the Samba daemon (smbd) listens for connections. 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 for SambaServer agent

```

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 for SambaServer agent

The sample configurations for this agent follow:

```

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

```

## Debug log levels for SambaServer agent

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 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 for SambaShare agent

SambaShare resources depend on the SambaServer, NetBios and Mount resources.

**Figure 4-5** Sample service group for a SambaShare resource



## Agent functions for SambaShare agent

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 connections with the particular samba share, removes its entry from the samba configuration file and reloads the configuration.

## State definitions for SambaShare agent

ONLINE	Indicates that the share is available.
OFFLINE	Indicates that the share is not available.
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 for SambaShare agent

**Table 4-8** Required attributes

Required attribute	Description
SambaServerRes	Name of the SambaServer resource. Type and dimension: string-scalar Example: "smb_res1"
ShareName	Name of the share resource as exported by samba. <b>Note:</b> This name can be different from the SambaShare resource name. Type and dimension: string-scalar Example: "share1"
ShareOptions	List of parameters for the share attributes. These parameters are specified as name=value pairs, with each pair separated by a semicolon (;). Type and dimension: string-scalar Example: "path=/shared; public=yes; writable=yes"

## Resource type definition for SambaShare agent

```

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 for SambaShare agent

```

SambaShare Samba_SambaShare3 (
  SambaServerRes = Samba_SambaServer
  ShareName = smbshare3
)

```

```
ShareOptions = "path=/smbshare3; public=yes; writable=yes"  
)
```

## Debug log levels for SambaShare agent

The SambaShare agent uses the following debug log levels:

DBG\_1, DBG\_3, DBG\_5

## NetBios agent

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

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:** The nmbd broadcasts the NetBIOS name, or the name by which the Samba server is known in the network.

---

## Dependencies for NetBios agent

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 for NetBios agent

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. Also verifies that the nmbd daemon is running by reading its pid file.
Clean	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, forcibly when necessary.

## State definitions for NetBios agent

ONLINE	Indicates that the specified NetBIOS name and 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 any 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 resource has become offline unexpectedly outside of VCS control.

## Attributes for NetBios agent

**Table 4-9** Required attributes

Required attribute	Description
NetBiosName	Name by which the Samba server is known in the network. Type and dimension: string-scalar Example: "samba_demon" <b>Note:</b> Samba has a limitation of 15 characters for NetBios names and aliases.
SambaServerRes	Name of the SambaServer resource. Type and dimension: string-scalar Example: "smb_res1"

**Table 4-10** Optional attributes

Optional attribute	Description
Interfaces	List of network interfaces on which Samba handles browsing. Type and dimension: string-vector Example: "172.29.9.24/16" <b>Note:</b> 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. Type and dimension: string-vector Example: { host1_samba, myname } <b>Note:</b> Samba has a limitation of 15 characters for NetBios names and aliases.
WinsSupport	If set to 1, this flag causes the agent to configure Samba as a WINS server. Type and dimension: integer-scalar Default: 0

**Table 4-10** Optional attributes (*continued*)

Optional attribute	Description
DomainMaster	<p>If set to 1, the agent sets Samba as Domain Master. Note that there can be only one domain master in a domain.</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 nonstandard configuration file name or path. If this attribute 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: "/var/run/nmbd.pid"</p>

## Resource type definition for NetBios agent

```

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 for NetBios agent

```

NetBios Samba_NetBios (
  SambaServerRes = Samba_SambaServer
  NetBiosName = samba_demon
  NetBiosAliases = { asamba_demon, samba127 }
)

```

```
WinsSupport = 1  
DomainMaster = 1  
)
```

## Debug log levels for NetBios agent

The NetBios agent uses the following debug log levels:

DBG\_1, DBG\_5

# Service and application agents

This chapter includes the following topics:

- [About the services and applications agents](#)
- [Apache HTTP server agent](#)
- [Application agent](#)
- [CoordPoint agent](#)
- [KVMGuest agent](#)
- [Process agent](#)
- [ProcessOnOnly agent](#)

## About the services and applications agents

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

## Apache HTTP server agent

The Apache HTTP server agent brings an Apache Server online, takes it offline, and monitors its processes. The Apache HTTP server agent consists of resource type declarations and agent scripts. You use the Apache HTTP server agent, in conjunction with other agents, to make an Apache HTTP 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, 2.0 and 7.0.0.0.

This agent can detect when an Apache HTTP 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 HTTP server notes”](#) on page 166.

## Dependencies

This type of resource depends on IP and Mount resources.

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



## Agent functions

Online

To start the Apache HTTP server, the agent:

- Executes the httpdDir/httpd program with the appropriate arguments if the httpdDir program specifies the full path of the directory in which the httpd binary file is located.
- Alternatively, if the httpdDir attribute specifies the full path of the Apache HTTP server binary file, the binary file is executed with appropriate arguments.

When you specify a file with the EnvFile attribute, the file is sourced before the agent executes the Apache HTTP server commands.



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, if httpdDir specifies the full path of the directory in which the httpd binary file is located.</li> <li>■ Alternatively, if the httpdDir attribute is used to specify the full path of the Apache HTTP server binary, the binary file is executed with appropriate arguments.</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 Apache HTTP server commands.</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.</p> <p>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 HTTP 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. Type and dimension: string-scalar Example: "/apache/server1/conf/httpd.conf"
httpdDir	Full path of the Apache HTTP server binary file or full path of the directory in which the httpd binary file is located. Type and dimension: string-scalar Example: "/apache/server1/bin"
PidFile	This attribute is required when you want to enable the detection of a graceful shutdown outside of VCS control. See <a href="#">Table 5-2</a> on page 162.

**Table 5-2** Optional attributes

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

Table 5-2 Optional attributes (*continued*)

Optional attribute	Description
EnableSSL	<p>If this attribute is set to 1 (true) the online agent function will add support for SSL, by including the option <code>-DSSL</code> in the start command.</p> <p>For example: <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>If this attribute is set to 0 (false) the agent excludes the SSL support.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: "1"</p>
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 <code>HostName</code> attribute.</p> <p><b>Note:</b> The <code>HostName</code> attribute is only required when the value of <code>SecondLevelMonitor</code> 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 <code>SecondLevelMonitor</code> 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 Apache HTTP server commands. With Apache 2.0, the file <code>ServerRoot/bin/envvars</code>, which is supplied in most Apache 2.0 distributions, is commonly used to set the environment before executing <code>httpd</code>. Specifying this attribute is optional. If <code>EnvFile</code> is specified, the shell for user 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 (*continued*)

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 Apache HTTP server. For the agent to detect the graceful shutdown of the Apache HTTP 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 <code>httpd</code> program. Refer to the <code>httpd</code> 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).</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: <code>"1"</code></p>

**Table 5-2** Optional attributes (*continued*)

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 Table</p>
ResLogLevel	<p>Controls the agent's logging detail for a specific instance of a resource. Values are</p> <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> <p>Type and dimension: string-scalar</p> <p>Default: INFO</p> <p>Example: "TRACE"</p>

**Table 5-3** Resource type attribute

Optional 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 = { "checkconffile.vfd" }
    static str ArgList[] = { ResLogLevel, State, IState, httpdDir,
        SharedObjDir, EnvFile, PidFile, HostName, Port, User,
        SecondLevelMonitor, SecondLevelTimeout, ConfigFile, EnableSSL,

```

```
DirectiveAfter, DirectiveBefore }
str ResLogLevel = INFO
str httpdDir
str SharedObjDir
str EnvFile
str PidFile
str HostName
int Port = 80
str User
int SecondLevelTimeout = 30
str ConfigFile = 0
str DirectiveAfter{}
str DirectiveBefore{}

static boolean IntentionalOffline = 0
)
```

## Apache HTTP server notes

The Apache Apache HTTP server has the following notes:

- See [“Tasks to perform before you use the Apache HTTP server agent”](#) on page 166.
- See [“About detecting application failure”](#) on page 167.
- See [“About bringing an Apache HTTP server online outside of VCS control”](#) on page 167.
- See [“About high Availability fire drill”](#) on page 168.

## Tasks to perform before you use the Apache HTTP 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 HTTP 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 HTTP server online outside of VCS control”](#) on page 167.
- 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.
- 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.

If the binary file `ab` is not found, Apache agent uses the `ab2` binary file for detail monitoring.

## About bringing an Apache HTTP server online outside of VCS control

When you bring an Apache HTTP 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 HTTP server outside of VCS control

- 1 Source the environment file if required.
- 2 Start the Apache HTTP 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
```

---

**Note:** You can specify the full path of a binary file without having `httpd` as part of `httpdDir` attribute.

For example:

---

```
/usr/sbin/apache2 -f /etc/httpd/conf/httpd.conf -k start
```

## 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 HTTP 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

A sample main.cf file follows:

```
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
        FsckOpt = "-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 can be on local storage or shared storage. 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*.

Prevention Of Concurrency Violation (ProPCV), a group level attribute, can be enabled through the Command Line Interface (CLI) to prevent an online resource on a node from coming online on another node, outside of VCS control, in the same cluster. This action prevents data corruption of resources. The attribute can only be set for a local failover type group. For more information about ProPCV action entry point, refer to the *Veritas Cluster Server Administrator's Guide*.

## High availability fire drill for Application agent

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 refer to the *Veritas Cluster Server Administrator's Guide*.

## Dependencies for Application agent

No fixed dependency exists for Application agent.

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 for Application agent

Online	<p>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.</p> <p>To bring the resource online, the agent function performs the command:</p> <pre>su [-] user -c command_to_online_resource</pre>
Offline	<p>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.</p> <p>To take the resource offline, the agent function performs the command:</p> <pre>su [-] user -c command_to_offline_resource</pre>

Monitor	<p>If you specify the <code>MonitorProgram</code> attribute, the agent executes the user defined <code>MonitorProgram</code> in the user-specified context. If you specify the <code>PidFiles</code> attribute, the routine verifies that the process ID that is found in each listed file is running. If you specify the <code>MonitorProcesses</code> attribute, the routine verifies that each listed process is running in the context of the user you specify.</p> <p>Use any combination among these attributes (<code>MonitorProgram</code>, <code>PidFiles</code>, or <code>MonitorProcesses</code>) to monitor the application.</p> <p>If any of the processes that are specified in either <code>PidFiles</code> or <code>MonitorProcesses</code> is determined not to be running, the monitor returns <code>OFFLINE</code>. If the process terminates ungracefully, the monitor returns <code>OFFLINE</code> and failover occurs.</p> <p>To monitor the resource, the agent function performs the command:</p> <pre>su [-] user -c <i>command_to_monitor_resource</i></pre>
<code>imf_init</code>	<p>Initializes the agent to interface with the asynchronous monitoring framework (AMF) kernel driver. This function runs when the agent starts up.</p>
<code>imf_getnotification</code>	<p>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.</p>
<code>imf_register</code>	<p>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 <code>PidFiles</code> and the <code>MonitorProcesses</code> attribute.</p>
Clean	<p>Terminates processes specified in <code>PidFiles</code> or <code>MonitorProcesses</code>. Ensures that only those processes (that are specified in the <code>MonitorProcesses</code> attribute) running with the user ID specified in the <code>User</code> attribute are killed. If the <code>CleanProgram</code> is defined, the agent executes the <code>CleanProgram</code>.</p> <p>To forcefully stop the resource, the agent function performs the command:</p> <pre>su [-] user -c <i>command_to_clean_resource</i></pre> <p>Note that the agent uses the <code>su -</code> option only when the attribute <code>UseSUDash</code> is enabled (1). The <code>UseSUDash</code> attribute is disabled (0) by default.</p>



Action	<p>The various functions of the action entry point are as follows:</p> <ul style="list-style-type: none"><li>■ <code>program.vfd</code> Checks the availability of the specified program and the execution permissions for the specified program.</li><li>■ <code>user.vfd</code> Checks the existence of the specified user on the host.</li><li>■ <code>cksum.vfd</code> Checks the existence of the same binary on all nodes.</li><li>■ <code>propcv</code> Invokes the AMF call with arguments to decide whether to allow or prevent processes from starting for an application resource, outside the VCS control, in the cluster.</li><li>■ <code>getcksum</code> Returns the checksum of the specified program</li></ul>
--------	--

## State definitions for Application agent

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 for Application agent

**Table 5-4** Required attributes for Linux

Required attribute	Description
StartProgram	<p>The executable which starts the application. The attribute specifies the complete path of the executable. Applicable command line arguments follow the name of the executable. The executable path and argument have spaces separating them. This executable can be on a local storage or on a shared storage.</p> <p>For example, if the attribute for StartProgram is  <code>/usr/sbin/vxnotify -g dg00 -m &gt;&gt; /var/log/vxnotify.log</code>            (and vxnotify is blocking command) set it like:  <code>/usr/sbin/vxnotify -g dg00 -m &gt;&gt; /var/log/vxnotify.log &amp;</code></p> <p><b>Note:</b> Do not use the opening and closing ({} ) brace symbols in this string.</p> <p>Type and dimension: string-scalar            Example: <code>"/usr/sbin/sample_app start"</code></p>
StopProgram	<p>The executable which stops the application. The Attribute specifies the complete path of the executable. Applicable command line arguments follow the name of the executable. The executable path and argument have spaces separating them. This executable can be on a local storage or on a shared storage.</p> <p><b>Note:</b> Do not use the opening and closing ({} ) brace symbols in this string.</p> <p>Type and dimension: string-scalar            Example: <code>"/usr/sbin/sample_app stop"</code></p>
At least one of the following attributes: <ul style="list-style-type: none"> <li>■ MonitorProcesses</li> <li>■ MonitorProgram</li> <li>■ PidFiles</li> </ul>	See <a href="#">Table 5-5</a> on page 179.

Table 5-5 Optional attributes for Linux

Optional attribute	Description
CleanProgram	<p>The executable which forcibly stops the application. The Attribute specifies the complete path of the executable. Applicable command line arguments follow the name of the executable. The executable path and argument have spaces separating them. This executable can be on a local storage or on a shared storage.</p> <p><b>Note:</b> Symantec recommends to have the CleanProgram on the local storage so that in case of loss of storage connectivity VCS can take appropriate action to stop the application.</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 which monitors the application. The Attribute specifies the complete path of the executable. Applicable command line arguments follow the name of the executable. The executable path and argument have spaces separating them. This executable can be on a local storage or on a shared storage.</p> <p>MonitorProgram can return the following VCSAgResState values: OFFLINE value is 100 or 1; ONLINE values range from 101 to 110 or 0 (depending on the confidence level); 110 equals confidence level of 100%. Any other value = UNKNOWN.</p> <p><b>Note:</b> Do not use the opening and closing ({} ) brace symbols in this string.</p> <p>If MonitorProgram is configured and not available, then resource state will be:</p> <ul style="list-style-type: none"> <li>■ OFFLINE if the resource was in OFFLINE state and not waiting for any action.</li> <li>■ UNKNOWN if the resource was in any other state or waiting for some action.</li> </ul> <p>Type and dimension: string-scalar</p> <p>Example: "/usr/sbin/sample_app_monitor all"</p>

**Table 5-5** Optional attributes for Linux (*continued*)

Optional attribute	Description
PidFiles	<p>A list of PID (process ID) files that contain the PID of the processes that you want monitored and cleaned. These are application generated files. Each PID file contains one monitored PID. Specify the complete path of each PID file in the list.</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 ID for running StartProgram, StopProgram, MonitorProgram, and CleanProgram. The processes that are specified in the MonitorProcesses list must run in the context of the specified user. Monitor checks the processes to make sure they run in this context.</p> <p>Type and dimension: string-scalar</p> <p><b>Note:</b> If configured user does not exist then the resource state will be UNKNOWN.</p> <p><b>Note:</b> If home directory is not set for a configured user then the resource state will be UNKNOWN.</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><b>Note:</b> Please make sure that the EnvFile adheres the default shell syntax of the configured user.</p> <p>Example: /home/username/envfile</p>

**Table 5-5** Optional attributes for Linux (*continued*)

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 StartProgram, the StopProgram, the MonitorProgram, or the CleanProgram 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 StartProgram, the StopProgram, the MonitorProgram or the CleanProgram agent functions.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>

## Resource type definition for Application agent

```

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

```

## Notes for Application agent

### 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 the following table:

**Table 5-6**

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.

---

**Note:** For a resource, if a PID file configured in the PidFiles attribute and a process configured in the MonitorProcesses attribute have the same process ID (PID), then the resource fails to register to IMF.

---

### Requirement for programs

The programs specified in StartProgram, StopProgram, MonitorProgram, CleanProgram should not continuously write to STDOUT or STDERR. If required, please redirect STDOUT and STDERR to some file.

## Requirement for default profile

The default profile of configured user should not have any blocking command such as `bash` or any other command such as `exec` that changes the behavior of the shell. This may lead to unexpected behavior.

## Sample configurations for Application agent

The sample configurations for this agent follow:

### Configuration 1 for Application agent

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 for Application agent

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 for Application agent

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 `vxfsenadm` 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:

See [“Notes for the CoordPoint agent”](#) on page 186.

## Dependencies for CoordPoint agent

No dependencies exist for the CoordPoint resource.



## Agent functions for CoordPoint agent

Monitor	<p>Enables the CoordPoint agent to validate the node registrations in the coordination points and confirms that the coordination points are accessible.</p> <p>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 <code>hastatus -sum</code>.</p> <p>The CoordPoint agent also performs I/O fencing reporting activities.</p> <p>See <a href="#">“CoordPoint agent I/O fencing reporting activities”</a> on page 186.</p>
---------	---

## State definitions for CoordPoint agent

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 for CoordPoint agent

**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 FAULTED.</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 FAULTED 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</p> <p>Default: "0"</p>

## Resource type definition for CoordPoint agent

```
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 for CoordPoint agent

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 for CoordPoint agent

The CoordPoint agent uses the following debug log levels:

DBG\_1, DBG\_10

## KVMGuest agent

The KVMGuest agent monitors a Linux Kernel-based Virtual Machine (KVM guest) and brings it online and takes it offline. The KVMGuest agent uses virsh commands. You can use this agent to make a KVM guest highly available and to monitor it.

Veritas Cluster Server supports KVM guests created on Red Hat Enterprise Linux 6 and update 1.


For important information on this agent, refer to:

See [“Notes for KVMGuest agent”](#) on page 192.

## Dependencies for KVMGuest agent

The KVMGuest resource depends on the Mount resource. The KVMGuest resource can depend on CFSMount too.

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



## Agent functions for KVMGuest agent

Online	Uses the <code>virsh start</code> command to start the KVM guest. When the KVM guest is configured, use the <code>virsh define</code> command to define the KVM guest while bringing it online.
Offline	Uses the <code>virsh shutdown</code> command to shutdown the KVM guest. If the <code>SyncDir</code> attribute is configured to synchronize the KVM guest configuration file, then the configuration file is copied to the <code>SyncDir</code> attribute.
Monitor	Uses the <code>virsh dominfo</code> command to determine the status of the KVM guest. This function also checks the status of the KVM guest process.
Clean	Uses the <code>virsh destroy</code> command to forcefully shutdown the KVM guest. If the <code>SyncDir</code> attribute is configured to synchronize the KVM guest configuration file, then the configuration file is copied to the <code>SyncDir</code> attribute.
Open	Starts the <code>vcskvmgustedd</code> process which is used to detect the guest migration.
Close	Stops the <code>vcskvmgustedd</code> process
Shutdown	Stops the <code>vcskvmgustedd</code> process
attr_changed	checks the value of <code>DelayAfterGuestOnline</code> and <code>DelayAfterGuestOffline</code> attribute. If it is less than the default value then the warning message is logged.
Action	<code>guestmigrated</code>  <code>vcskvmgustedd</code> process detects the guest migration and executes this action entry point to create the migration state file and initiate the resource monitor using the <code>hares -probe</code> command.

## State definitions for KVMGuest agent

ONLINE	Indicates that the KVM guest is running.
OFFLINE	Indicates that the KVM guest has stopped.
FAULTED	Indicates that the KVM guest has failed to start or unexpectedly stopped.
UNKNOWN	Indicates that the problem exists with the configuration or the ability to monitor the resource.
INTENTIONAL OFFLINE	Indicates that the KVM guest is either migrated to another physical host or the KVM guest is intentionally paused by the administrator.

## Attributes for KVMGuest agent

**Table 5-8** Required attribute

Required attribute	Description
GuestName	The name of the KVM guest created using the KVM hypervisor. Type and dimension: string-scalar Example: kvmguest1

**Table 5-9** Optional attribute

Optional attribute	Description
DelayAfterGuestOnline	Defines the maximum time in seconds that the KVM guest takes to start. You can modify this attribute as required. Type and dimension: integer-scalar Default: 5 Example: 10

**Table 5-9** Optional attribute (*continued*)

Optional attribute	Description
DelayAfterGuestOffline	<p>Defines the maximum time in seconds that the KVM guest takes to shut down. You can modify this attribute as required.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 30</p> <p>Example: 40</p>
SyncDir	<p>Specifies the absolute path of a directory used to synchronize the default configuration directory /etc/libvirt/qemu/ on all the cluster nodes. If this attribute is configured, the online entry point will use the guest configuration file in the specified path to define the guest. This directory should be on the shared storage.</p> <p>Type and dimension: string-scalar</p> <p>Example: /kvmguest_syncdir</p>
GuestConfigFilePath	<p>Specifies the absolute path of the guest configuration file. You can use this attribute as an alternative to SyncDir. If you configure this attribute, ensure that the guest configuration file is available on all the cluster nodes and on same path. VCS uses this path to define the guest on each node. If both SyncDir and GuestConfigFilePath are configured, then SyncDir will be preferred over GuestConfigFilePath.</p> <p>Type and dimension: string-scalar</p> <p>Example: /kvmguest/kvmguest1.xml</p>

**Table 5-10** Internal attributes

Internal attribute	Description
IntentionalOffline	<p>For information on how to use the IntentionalOffline resource type attribute, refer to the <i>Veritas Cluster Server Administrator's Guide</i>.</p>
CEInfo	<p>This is an internal attribute.</p>

## Resource type definition for KVMGuest agent

```
type KVMGuest (
    static keylist SupportedActions = { "guestmigrated" }
    static str ArgList[] = { GuestName, DelayAfterGuestOnline,
        DelayAfterGuestOffline, SyncDir, GuestConfigFilePath, CEInfo }
    str CEInfo{} = { Enabled=0, CESystem=NONE, FaultOnHBLoss=1 }
    static int IntentionalOffline = 1
    str GuestName
    int DelayAfterGuestOnline = 5
    int DelayAfterGuestOffline = 30
    str SyncDir
    str GuestConfigFilePath
)
```

## Notes for KVMGuest agent

The KVMGuest agent has the following notes:

- [Support for guests created on Red Hat Enterprise Linux 6](#)
- [Storage and network configurations](#)
- [Guest live migration](#)

### Support for guests created on Red Hat Enterprise Linux 6

KVMGuest agent supports only those guests that can be created on Red Hat Enterprise Linux 6.

### Storage and network configurations

If you have specified customized network and storage configurations for any guest, make sure that same configurations are available on all the cluster nodes.

### Guest live migration

For migrating the guest from one node to another node, make sure that the guest image is available on both the nodes simultaneously at the same location. Also, KVM requires that the guest network and storage configuration should be identical on source and destination node. For details, refer to the *Redhat Enterprise Linux Virtualization Guide*.



## Sample configurations for KVMGuest agent

### Sample Configuration 1 for KVMGuest agent

In this example, the native LVM volumes are used to store the guest image.

```
group kvmtest1 (  
  SystemList = { north = 0, south = 1 }  
)  
KVMGuest res1 (  
  GuestName = kvmguest1  
  GuestConfigFilePath = "/kvmguest/kvmguest1.xml"  
  DelayAfterGuestOnline = 10  
  DelayAfterGuestOffline = 35  
)  
Mount mnt1 (  
  BlockDevice = "/dev/mapper/kvmvg-kvmvol"  
  MountPoint = "/kvmguest"  
  FSType = ext3  
  FsckOpt = "-y"  
  MountOpt = "rw"  
)  
LVMLogicalVolume lv1 (  
  VolumeGroup = kvmvg  
  LogicalVolume = kvmvol  
)  
LVMVolumeGroup vgl (  
  VolumeGroup = kvmvg  
)  
res1 requires mnt1  
mnt1 requires lv1  
lv1 requires vgl
```

### Sample Configuration 2 for KVMGuest agent

In this example, the native VxVM volumes are used to store the guest image.

```
group kvmtest2 (  
  SystemList = { north = 0, south = 1 }  
)  
KVMGuest res1 (  
  GuestName = kvmguest1  
  GuestConfigFilePath = "/kvmguest/kvmguest1.xml"  
  DelayAfterGuestOnline = 10
```

```
DelayAfterGuestOffline = 35
)
Mount mnt1 (
BlockDevice = "/dev/vx/dsk/kvmvg/kvmvol"
MountPoint = "/kvmguest"
FSType = vxfs
FsckOpt = "-y"
MountOpt = "rw"
)
Volume voll (
Volume = kvm_vol
DiskGroup = kvm_dg
)
DiskGroup dg1 (
DiskGroup = kvm_dg
)
res1 requires mnt1
mnt1 requires voll
voll requires dg1
```

### Sample Configuration 3 for KVMGuest agent

In this example, the native CVM-CFS is used to store the guest image.

```
group cvm (
SystemList = { kvmpm1 = 0, kvmpm2 = 1 }
AutoFailOver = 0
Parallel = 1
AutoStartList = { kvmpm1, kvmpm2 }
)

CFSSMount cfsmount (
MountPoint = "/cfsmount"
BlockDevice = "/dev/vx/dsk/cfsdg/cfsvol"
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
CVMClustName = kvmcfs
CVMNodeId = { kvmpm1 = 0, kvmpm2 = 1 }
CVMTransport = gab
CVMTimeout = 200
```

```
)

CVMVolDg cfsdg (
CVMDiskGroup = cfsdg
CVMVolume = { cfsvol }
CVMActivation = sw
)

CMMVxconfigd cvm_vxconfigd (
Critical = 0
CMMVxconfigdArgs = { syslog }
)

cfsmount requires cfsdg
cfsmount requires cvm_clus
cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus
```

## Debug log levels for KVMGuest agent

The DiskGroup agent uses the following debug log levels:

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

## 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 does an exact match of configured PathName and Arguments with the processes in the process table. To clarify, the processes that `ps` command lists should have exact match of PathName and Arguments for all the configured 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 for Process agent

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 refer to the *Veritas Cluster Server Administrator's Guide*.

## Dependencies for Process agent

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

**Figure 5-4** Sample service group for a Process resource



## Agent functions for Process agent

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.

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

## State definitions for Process agent

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 along with the arguments.
OFFLINE	Indicates that the specified process is not running.
FAULTED	Indicates that the process has terminated unexpectedly.
UNKNOWN	Indicates that the agent can not determine the state of the process.

## Attributes for Process agent

**Table 5-11** Required attribute for Linux

Required attribute	Description
PathName	<p>Absolute path 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 5-12** Optional attributes for Linux

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>
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>

**Table 5-12** Optional attributes for Linux (*continued*)

Optional attribute	Description
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 for Process agent

```

type Process (
    static keylist SupportedActions = { "program.vfd", getcksum }
    static str ArgList[] = { PathName, Arguments, UserName,
        Priority, PidFile }
    static int IMF{} = { Mode = 3, MonitorFreq = 5,
        RegisterRetryLimit = 3 }
    str PathName
    str Arguments
    str UserName = root
    str Priority = 10
    str PidFile
)

```

## Usage notes for Process agent

The Process agent has the following notes:

- [Requirement for programs](#)

### Requirement for programs

The programs specified in PathName should not continuously write to STDOUT or STDERR. If required, please redirect STDOUT and STDERR to some file.

## Sample configurations for Process agent

### Configuration for Process agent

Configuration for Linux follows:

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 for Process agent

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 or to monitor it. 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 along with the arguments.
FAULTED	Indicates that the process has unexpectedly terminated.
UNKNOWN	Indicates that the agent can not determine the state of the process.

## Attributes

**Table 5-13** Required attributes for Linux

Required attribute	Description
PathName	<p>Defines absolute path 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 5-14** Optional attributes for Linux

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>

**Table 5-14** Optional attributes for Linux (*continued*)

Optional attribute	Description
IgnoreArgs	<p>A flag that indicates whether monitor ignores the argument list.</p> <ul style="list-style-type: none"> <li>■ If the value is 0, it checks the process pathname and argument list.</li> <li>■ If the value is 1, it only checks for the executable pathname and ignores the rest of the argument list.</li> </ul> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>
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
    )
```

## ProcessOnOnly agent usage notes

The ProcessOnOnly agent has the following notes:

- [Requirement for programs](#)

### Requirement for programs

The programs specified in PathName should not continuously write to STDOUT or STDERR. If required, please redirect STDOUT and STDERR to some other file.

## 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 includes the following topics:

- [About the infrastructure and support agents](#)
- [NotifierMngr agent](#)
- [Proxy agent](#)
- [Phantom agent](#)
- [RemoteGroup agent](#)

## 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.

Refer to the *Admin 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 <code>notifier</code> 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 <code>notifier</code> process.
Clean	Sends <code>SIGKILL</code> .

## State definitions

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

## Attributes

**Table 6-1** Required attributes for Linux

Required attribute	Description
SnmpConsoles	<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>SnmpConsoles is a required attribute if SntpServer is not specified; otherwise, SnmpConsoles is an optional attribute. Specify both SnmpConsoles and SntpServer if desired.</p> <p>Type and dimension: string-association</p> <p>Example:  "172.29.10.89" = Error, "172.29.10.56" = Information</p>
SntpServer	<p>Specifies the machine name of the SMTP server.</p> <p>SntpServer is a required attribute if SnmpConsoles is not specified; otherwise, SntpServer is an optional attribute. You can specify both SntpServer and SnmpConsoles if desired.</p> <p>Type and dimension: string-scalar</p> <p>Example: "smtp.example.com"</p>

**Table 6-2** Optional attributes for Linux

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>
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>

**Table 6-2** Optional attributes for Linux (*continued*)

Optional attribute	Description
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><b>Note:</b> 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>
SmtpFromPath	<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>
SmtpRecipients	<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><b>Note:</b> SmtpRecipients is a required attribute if you specify SmtpServer.</p> <p>Type and dimension: string-association</p> <p>Example:</p> <pre>"james@example.com" = SevereError, "admin@example.com" = Warning</pre>
SmtpReturnPath	<p>Set to a valid email address, if you want the notifier to use a custom email address in the Return-Path: &lt;&gt; field.</p> <p>If the mail server specified in SmtpServer does not support SMTP VRFY command, then you need to set the SmtpVrfyOff to 1 in order for the SmtpReturnPath value to take effect.</p> <p>Type and dimension: string-scalar</p> <p>Example: "usera@example.com"</p>

**Table 6-2** Optional attributes for Linux (*continued*)

Optional attribute	Description
SmtplibServerTimeout	This attribute represents the time in seconds notifier waits for a response from the mail server for the SMTP commands it has sent to the mail server. This value can be increased if you notice that the mail server is taking a longer duration to reply back to the SMTP commands sent by notifier.  Type and dimension: integer-scalar Default: 10
SmtplibServerVrfyOff	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 SmtplibServer attribute while sending emails.  Type and dimension: boolean-scalar Default: 0
SnmplibCommunity	Specifies the community ID for the SNMP manager.  Type and dimension: string-scalar Default: public
SnmplibTrapPort	Port on the SNMP console machine where SNMP traps are sent.  If you specify more than one SNMP console, all consoles use this value.  Type and dimension: string-scalar Default: 162

## Resource type definition

```

type NotifierMngr (
  static int RestartLimit = 3
  static str ArgList[] = { EngineListeningPort, MessagesQueue,
  NotifierListeningPort, NotifierSourceIP, SnmplibTrapPort,
  SnmplibCommunity, SnmplibConsoles, SmtplibServer, SmtplibServerVrfyOff,
  SmtplibServerTimeout, SmtplibReturnPath, SmtplibFromPath, SmtplibRecipients
}
  int EngineListeningPort = 14141
  int MessagesQueue = 30
  int NotifierListeningPort = 14144

```



```
str NotifierSourceIP
int SnmpdTrapPort = 162
str SnmpCommunity = public
str SnmpConsoles{}
str SmtperServer
boolean SmtperServerVrfyOff = 0
int SmtperServerTimeout = 10
str SmtperReturnPath
str SmtperFromPath
str SmtperRecipients{}
str NotifierSourceIP
)
```

## Sample configuration

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

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

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

---

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

---

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

## Configuration

Configuration for Linux follows:

```
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 either in a failover or a 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: "nic1"</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

```

Proxy proxy1 (
    TargetResName = "nic1"
)

```

## Configuration 2

The proxy resource mirrors the state of the resource nic2 on sysa.

```
Proxy proxy1 (
    TargetResName = "nic2"
    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 group cannot be in the same group.

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

MultiNICA mnic (
    Device@sysa = { eth0 = "192.123.8.41", eth3 =
        "192.123.8.42" }
    Device@sysb = { 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 = "255.255.248.0"  
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 agent enables VCS to determine the status of parallel service groups that do not include OnOff resources, which are resources that VCS can start and stop. Without the "dummy" resource provided by this agent, VCS cannot assess the status of groups that only contain None (Persistent) and OnOnly resources because the state of these resources is not considered in the process of determining whether a group is online. Refer to the *VCS Administrator's Guide* for information on categories of service groups and resources.

Do not use the Phantom resource in failover service groups.

Also, the Phantom resource should not be used in service groups that don't contain any resources.

---

**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 (  
)
```

## 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 refer to 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. For more information: See <a href="#">Table 6-5</a> on page 218.
Offline	Takes the remote service group offline. For more information: See <a href="#">Table 6-5</a> on page 218.
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. For more information: See <a href="#">Table 6-5</a> on page 218.
Clean	If the RemoteGroup resource faults, the Clean function takes the remote service group offline. For more information: See <a href="#">Table 6-5</a> on page 218.

## 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.
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.
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 <a href="#">Table 6-6</a> on page 221.</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: "DBGrp"</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"><li>■ ANY The RemoteGroup resource goes online if the remote service group is online on any node in the remote cluster.</li><li>■ <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.</li></ul> <p>Type and dimension: string-scalar</p> <p>Example: "vcssys1" or "ANY"</p>

**Table 6-5** Required attributes (*continued*)

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"> <li>■ OnOff           <p>The RemoteGroup resource brings the remote service group online or takes it offline.</p> <p>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> </li> <li>■ MonitorOnly           <p>The RemoteGroup resource only monitors the state of the remote service group. The RemoteGroup resource cannot online or offline the remote service group.</p> <p>Make sure that you bring the remote service group online before you online the RemoteGroup resource.</p> </li> <li>■ OnlineOnly           <p>The RemoteGroup resource only brings the remote service group online. The RemoteGroup resource cannot take the remote service group offline.</p> <p>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> </li> </ul> <p>Type and dimension: string-scalar</p>

**Table 6-5** Required attributes (*continued*)

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"> <li>■ 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.</li> <li>■ 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.</li> </ul> <p>Type and dimension: string-scalar</p> <p>Example:</p> <ul style="list-style-type: none"> <li>■ For a cluster without the Symantec Product Authentication Service: "johnsmith"</li> <li>■ For a secure remote cluster: "foobar@example.com"</li> </ul>
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><b>Note:</b> 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 (*continued*)

Optional attribute	Description
ReturnIntOffline	<p>Select one of the following values for RemoteGroup to return IntentionalOffline:</p> <ul style="list-style-type: none"> <li>■ RemotePartial—Indicates that the RemoteGroup resource returns an IntentionalOffline if the remote service group is in an ONLINE PARTIAL state.</li> <li>■ RemoteOffline—Indicates that the RemoteGroup resource returns an IntentionalOffline if the remote service group is in an OFFLINE state.</li> <li>■ RemoteFaulted—Indicates that the RemoteGroup resource returns an IntentionalOffline if the remote service group is OFFLINE FAULTED.</li> </ul> <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. 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>
ToleranceLimit MonitorInterval	<p>If you expect the RemoteGroup agent to tolerate sudden offlines of the remote service group, then modify the ToleranceLimit attribute.</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 (*continued*)

Type level attributes	Description
ExternalStateChange	<p>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.</p> <p>See the <i>Veritas Cluster Server Administrator's Guide</i> for more information about these attributes.</p>

## 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 includes the following topics:

- [About the testing agents](#)
- [ElifNone agent](#)
- [FileNone agent](#)
- [FileOnOff agent](#)
- [FileOnOnly agent](#)

## 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 for ElifNone agent

No dependencies exist for the ElifNone resource.



**Figure 7-1** Sample service group that includes an ElifNone resource



## Agent function for ElifNone agent

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

## State definitions for ElifNone agent

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 for ElifNone agent

**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 for ElifNone agent

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

## Sample configuration for ElifNone agent

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

## Debug log levels for ElifNone agent

The ElifNone agent uses the following debug log levels:  
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.

See ["Dependencies for FileNone agent"](#) on page 226.

See ["Agent functions for FileNone agent"](#) on page 227.

See ["State definitions for FileNone agent"](#) on page 227.

See ["Attributes for FileNone agent"](#) on page 228.

See ["Resource type definition for FileNone agent"](#) on page 228.


See ["Sample configuration for FileNone agent"](#) on page 228.

See ["Debug log levels for FileNone agent"](#) on page 228.

## Dependencies for FileNone agent

No dependencies exist for the FileNone resource.

**Figure 7-2** Sample service group that includes an FileNone resource



See [“FileNone agent”](#) on page 226.

## Agent functions for FileNone agent

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

See [“FileNone agent”](#) on page 226.

## State definitions for FileNone agent

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.

See [“FileNone agent”](#) on page 226.

## Attributes for FileNone agent

**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"

See [“FileNone agent”](#) on page 226.

## Resource type definition for FileNone agent

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

See [“FileNone agent”](#) on page 226.

## Sample configuration for FileNone agent

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

See [“FileNone agent”](#) on page 226.

## Debug log levels for FileNone agent

The FileNone agent uses the following debug log levels:

DBG\_4, DBG\_5

See [“FileNone agent”](#) on page 226.

## FileOnOff agent


The FileOnOff agent creates, removes, and monitors a file.

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 for FileOnOff agent

No dependencies exist for the FileOnOff resource.

**Figure 7-3** Sample service group that includes a FileOnOff resource



## Agent functions for FileOnOff agent

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	Removes the specified file forcibly when necessary.

## State definitions for FileOnOff agent

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 for FileOnOff agent

**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 for FileOnOff agent

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

## Sample configuration for FileOnOff agent

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

## Debug log levels for FileOnOff agent

The FileOnOff agent uses the following debug log levels:

DBG\_1, DBG\_4, DBG\_5


## FileOnOnly agent

The FileOnOnly agent creates and monitors a file.

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 for FileOnOnly agent

No dependencies exist for the FileOnOnly resource.

**Figure 7-4** Sample service group that includes a FileOnOnly resource

## Agent functions for FileOnOnly agent

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 for FileOnOnly agent

The state definitions for this agent follow:

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 for FileOnOnly agent

**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 for FileOnOnly agent

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

## Sample configuration for FileOnOnly agent

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

## Debug log levels for FileOnOnly agent

The FileOnOnly agent uses the following debug log levels:

DBG\_1, DBG\_4, DBG\_5



# Replication agents

This chapter includes the following topics:

- [About the replication agents](#)
- [RVG agent](#)
- [RVGPrimary agent](#)
- [RVGSnapshot](#)
- [RVGShared agent](#)
- [RVGLogowner agent](#)
- [RVGSharedPri agent](#)

## About the replication agents

Use the replication agents to provide high availability for VVR resources.

Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for information on configuring the Replication agents for high availability.

## RVG agent

Brings the RVG online, monitors read and write access to the RVG, and takes the RVG offline. This is a failover resource. The RVG agent enables replication between clusters. It manages the Primary VVR node in one cluster and the Secondary VVR node in another cluster. Each node can be failed over in its respective cluster. In this way, replication is made highly available.

The RVG agent manages the state of the RVG during local failovers. The RVGPrimary agent manages the role of the RVG during a wide area failover.

Using a VCS global cluster enables you to fail over the Primary role from a Primary VVR node to a Secondary VVR node.

The RVG agent includes the following key features:

- Removes potential single points of failure by enabling Primary and Secondary VVR nodes to be clustered.
- Enables you to bring a service group online to start VCS-managed applications that use VVR.
- Continues replication after a node in a cluster fails without losing updates.
- Ensures that VVR can be added to any VCS cluster by including the RVG resource type definitions.

An example configuration file for this agent that can be used as a guide when creating your configuration is located at:

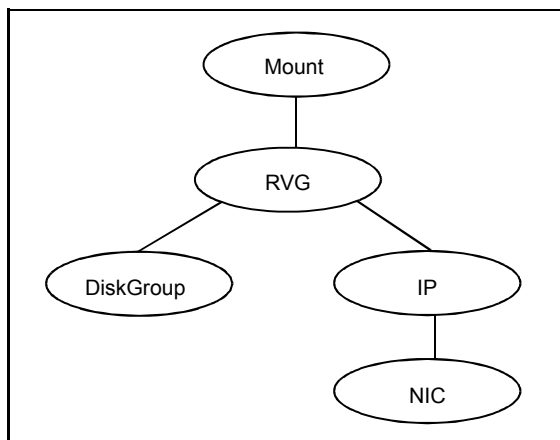
```
/etc/VRTSvcs/conf/sample_vvr/RVG
```

## Dependencies for RVG agent

The RVG resource represents the RVG (Replicated Volume Group) in the RDS (Replicated Data Set). The RVG resource is dependent on the DiskGroup resource. The RVG resource is also dependent on the IP resources that it uses for replication.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information on dependencies.

**Figure 8-1** Sample service group for an RVG resource



## Agent functions for RVG agent

The RVG agent has the following agent functions:

Online	Verifies whether the DiskGroup agent has recovered the RVG. If not, recovers and starts the data volumes and the Storage Replicator Log (SRL), recovers the RVG, recovers all RLINKs in the RVG, and then starts the RVG.
Offline	Stops the RVG.
Monitor	Monitors the state of the RVG using the <code>vxprint</code> command.  The RVG resource monitors an RVG for local access only. It does not monitor replication.
Clean	Stops the RVG.
Info	The info entry point displays information about the replication status of a RDS.

## State definitions for RVG agent

The RVG agent has the following state definitions:

ONLINE	Indicates that the RVG is in <code>ENABLED/ACTIVE</code> state.
OFFLINE	Indicates that the RVG is in <code>DISABLED/CLEAN</code> state.
FAULTED	The RVG resource fails if the RVG is not in the <code>ENABLED/ACTIVE</code> state.

## Attributes for RVG agent

**Table 8-1** Required attributes

Required attributes	Description
RVG	The name of the RVG being monitored. Type and dimension: string-scalar Example: "hr_rvg"
DiskGroup	The disk group that this RVG is associated with. Type and dimension: string-scalar Example: "hrbg"

**Table 8-1** Required attributes (*continued*)

Required attributes	Description
StorageDG	The name of the bunker disk group. Type and dimension: string-scalar Example: "hr_bdg"
StorageRVG	The name of the bunker RVG. Type and dimension: string-scalar Example: "hr_brvg"
StorageHostIds	A space-separated list of the hostids of each node in the bunker cluster. Type and dimension: string-keylist Example: "bunker_host"

## Resource type definitions for RVG agent

The RVG agent resource type definition follows.

```
type RVG (  
    static int NumThreads = 1  
    static str ArgList[] = { RVG, DiskGroup }  
    str RVG  
    str DiskGroup  
    str StorageRVG  
    str StorageDG  
    str StorageHostIds  
)
```

## Sample configurations for RVG agent

```
RVG rvg (  
    RVG = ApplicationRVG  
    DiskGroup = vvrldg  
    StorageRVG = ApplicationRVG  
    StorageDG = vvrldg  
    StorageHostIds = "bunker_host"  
)
```

## RVGPrimary agent

The RVGPrimary agent enables migration and takeover of a VVR Replicated Volume Group (RVG) in a VCS environment. Bringing a resource of type RVGPrimary online causes the RVG on the local host to become a primary.

The agent is useful when hosts in both the primary and secondary side are clustered, in particular a VCS replicated data cluster or a VCS global cluster, to completely automate the availability of writable replicated disks to a VCS-managed application.

The RVGPrimary agent includes the following features:

- Removes the manual steps of migrating a VVR primary and secondary roles when failing over applications across a wide area.
- Minimizes the need for resynchronizing replicated volumes by attempting a migration before attempting a hard takeover.
- Waits for the two sides of a replicated data set to become completely synchronized before migrating roles.
- Supports an automatic fast failback resynchronization of a downed primary if it later returns after a takeover.
- Allows you to distinguish the Primary site after network failure or disaster
- Supports the ability to choose the Primary site after a site failure or network disruption is corrected.
- After a successful migration or takeover of a Secondary RVG, the RVGPrimary agent ensures to automatically start the replication from the new Primary to any additional Secondary(s) that exists in the RDS.
- Before a takeover, the RVGPrimary agent synchronizes the Secondary site with any bunker associated with the Primary site, when the Primary site is not available.

Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for information on configuring the Replication agents for high availability.

A sample configuration file for this agent that you can use as a guide to create the configuration is located at `/etc/VRTSvc/conf/sample_vvr/RVGPrimary`.

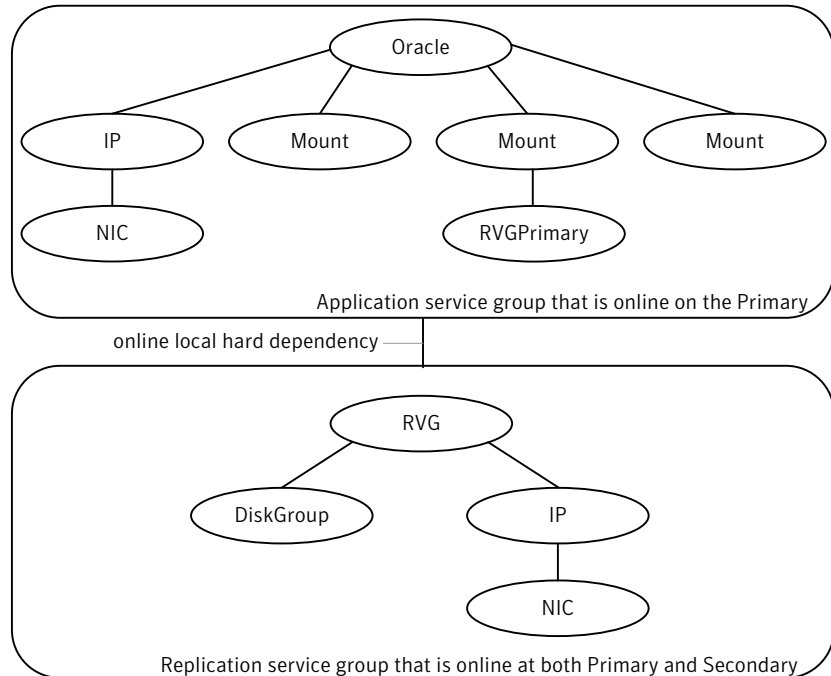
## Dependencies for RVGPrimary agent

You usually use the RVGPrimary agent in conjunction with the RVG agent in two groups with an online local hard group dependency. The parent group contains the resources that manage the actual application and file systems and as the

RVGPrimary resource. The child group contains the resources managing the storage infrastructure, which include the RVG and DiskGroup type resources.

Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for information about the setup of a VVR environment using the RVGPrimary agent.

**Figure 8-2** Sample service group for an RVGPrimary resource



## Agent functions for RVGPrimary agent

The RVGPrimary agent has the following agent functions:

- Online**                      Determines the current role of the RVG. If the role is Secondary it attempts a migration. It waits for any outstanding writes from the original Primary. If the original Primary is down, it attempts a takeover. You can configure the RVGPrimary agent so that, before a takeover, the agent synchronizes the Secondary site with any bunker associated with the Primary site, when the Primary site is not available. If the RVG is a Primary, it performs no actions and goes online.
- Offline**                      Performs no actions.

Monitor	Performs no actions. The RVG agents monitors the actual RVG.
Clean	Performs no actions.
fbsync	This is an action entry point.  It resynchronizes the original Primary with the new Primary that has taken over with fast-failback, after the original Primary had become unavailable. This needs to be executed when the original Primary becomes available and starts acting as a Secondary.
ElectPrimary	This is an action entry point.  It can be executed to retain the specified RVG as the Primary in a Primary-Primary configuration.  For more details, refer to the <i>Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide</i> .

## State definitions for RVGPrimary agent

The RVGPrimary agent has the following state definitions:

ONLINE	Indicates that the role of the RVG is Primary.
FAULTED	The RVG agents monitors the actual RVG. Accidental migration of a VVR Primary outside of VCS causes other resources to fault immediately, such as Mount. No special monitoring by this agent is necessary.

## Attributes for RVGPrimary agent

**Table 8-2** Required attributes

Required attributes	Description
RvgResourceName	The name of the RVG resource type that this agent promotes. The name RVG resource type which has been configured using the RVG agent.  Type and dimension: string-scalar
AutoTakeover	A flag to indicate whether the agent should perform a takeover on online if the original Primary is down.  AutoTakeover and AutoResync are mutually exclusive attributes.  When AutoTakeover=0, the primary-elect feature is not applicable; therefore, it is not supported.  Type and dimension: integer-scalar

**Table 8-2** Required attributes (*continued*)

Required attributes	Description
AutoResync	<p data-bbox="588 326 1220 409">Indicates whether the agent should attempt to automatically perform a fast-failback resynchronization of the original Primary after a takeover and after the original Primary returns.</p> <p data-bbox="588 427 1056 447">You can use the following values for this attribute:</p> <ul data-bbox="588 470 1220 791" style="list-style-type: none"><li data-bbox="588 470 1220 553">■ 0—instructs the agent to not attempt to perform a fast-failback resynchronization of the original Primary after a takeover and after the original Primary returns.</li><li data-bbox="588 562 1220 645">■ 1—instructs the agent to attempt to automatically perform a fast-failback resynchronization of the original Primary after a takeover and after the original Primary returns.</li><li data-bbox="588 654 1220 791">■ 2—instructs the agent to use the primary-elect feature. The agent does not attempt to perform a fast-failback resynchronization of the original Primary after a takeover and after the original Primary returns. The RVGPrimary agent also creates space-optimized snapshots for all the data volumes in the RVG resource.</li></ul> <p data-bbox="588 800 1220 942">If you set the AutoResync attribute to 2 (to enable the primary-elect feature) the value of the BunkerSyncTimeOut attribute must be zero to disable the automated bunker replay feature. You cannot use the automated bunker replay feature and the primary-elect feature in the same environment.</p> <p data-bbox="588 960 1197 980">AutoTakeover and AutoResync are mutually exclusive attributes.</p> <p data-bbox="588 1003 1220 1055">When AutoTakeover=0, the primary-elect feature is not applicable; therefore, it is not supported.</p> <p data-bbox="588 1078 921 1098">Type and dimension: integer-scalar</p>



**Table 8-2** Required attributes (*continued*)

Required attributes	Description
BunkerSyncTimeOut	<p>The value for the BunkerSyncTimeOut attribute determines if you want the bunker to perform a replay or not. You set the value in seconds for the time that you want to allot for the replay.</p> <p>Use one of the following values for the BunkerSyncTimeOut attribute:</p> <ul style="list-style-type: none"><li>■ If you do not use a value for this attribute (the default null value), the RVGPrimary agent considers it an infinite timeout value. The agent replays all the writes on the Bunker Replicator Log to the Secondary. Only after the agent sends all the writes, VCS performs the takeover on the Secondary.</li><li>■ If you set the value for this attribute to 0, you disable bunker replay for the agent. The RVGPrimary agent immediately performs a takeover on the Secondary. The agent does not send pending writes from the Bunker to the Secondary.</li><li>■ If you set the value to a number of seconds, then the RVGPrimary agent sends writes for that amount of time to the Secondary. After the agent meets the time limit, it performs the takeover on the Secondary. The bunker replay time in this case is equal to the value in seconds. You can set this value dynamically.</li></ul> <p>The RVGPrimary agent's OnlineTimeout and OnlineRetryLimit attribute values determine the available time for an RVGPrimary resource to complete its online operation.</p> <p>Use the following formula to get the Time Available for Online to Complete (TAOC):</p> $\text{TAOC} = (\text{OnlineTimeout} + (\text{OnlineRetryLimit} * \text{OnlineTimeout}))$

**Table 8-2** Required attributes (*continued*)

Required attributes	Description
BunkerSyncTimeOut (cont.)	<p>When you set the BunkerSyncTimeOut value in seconds, the value of TAOC for the RVGPrimary agent should be greater than the desired BunkerSyncTimeOut value. Using a TAOC value that is greater than BunkerSyncTimeOut value ensures that the bunker replay and the RVG takeover can complete in the allotted time for that particular online operation. If the TAOC is smaller than BunkerSyncTimeOut value and the bunker replay does not complete within the allotted time for the online process, the resource faults. If the resource faults, clear the fault. Try the online operation again if the resource has not failed over to other cluster node in the configuration.</p> <p>If you increase the value of the BunkerSyncTimeOut attribute, you need to increase the value of the OnlineTimeout or OnlineRetryLimit attribute so that TAOC remain greater than changed value. This is to ensure to have bunker replay completed within allotted time for online.</p> <p>If the value of the AutoResync attribute is 2, you must set the value of the BunkerSyncTimeOut attribute to 0 (to disable automated bunker replay).</p> <p>Type and dimension: string-scalar</p> <p>Default value: ""</p>

**Table 8-3** Internal attribute

Internal attribute	Description
BunkerSyncElapsedTime	<p>For internal use only, do not modify. This value in seconds signifies the amount of time that a Secondary RVG has waited for synchronization from the bunker host to complete.</p> <p>Type and dimension: integer-scalar</p>

## Resource type definitions for RVGPrimary agent

The RVGPrimary resource type definition follows.

```

type RVGPrimary (
    static keylist SupportedActions = { fbsync, electprimary }
    static int NumThreads = 1
    static int OnlineRetryLimit = 1
    static str ArgList[] = { RvgResourceName, "RvgResourceName:RVG",

```

```

    "RvgResourceName:DiskGroup", AutoTakeover, AutoResync,
    BunkerSyncTimeout, BunkerSyncElapsedTime }
    str RvgResourceName
    int AutoTakeover = 1
    int AutoResync = 0
    str BunkerSyncTimeout
    int BunkerSyncElapsedTime = 0
)

```

## Sample configurations for RVGPrimary agent

```

RVGPrimary rvg-pri (
    RvgResourceName = rvgRes
)

```

## RVGSnapshot

For a fire drill, creates and destroys a transactionally consistent space-optimized snapshot of all volumes in a VVR secondary replicated data set. The RVGSnapshot agent takes space-optimized snapshots on a secondary RVG. These snapshots can be mounted and written to without affecting the actual replicated data, which means that the space-optimized snapshot can be an effective tool for scheduling a “fire drill” to confirm that a wide-area failover is possible. By combining this agent with the VCS Mount agent, the CFMount agent, and VCS agents that manage the application being replicated, you can create a special fire drill service group. You can bring this service group online and take it offline at regularly scheduled intervals to verify that the disaster recovery environment is robust.

In addition to the agent itself, a text-based wizard `/opt/VRTSvcs/bin/fdsetup` that prepares the VVR and VCS infrastructure for a fire drill and a script `/opt/VRTSvcs/bin/fdsched` that runs the fire drill and consolidates the results are also included.

Complete details are in the *Veritas Cluster Server Administrator's Guide*.

The RVGSnapshot agent includes the following key features:

- Automates the process of creating a space-optimized snapshot on a VVR secondary that can be mounted to simulate a wide-area failover without affecting the production application.
- Includes a wizard to effectively set up and schedule fire drills that are completely managed by VCS.

---


**Note:** The RVGSnapshot agent does not support Volume Sets.

---

## Dependencies for RVGSnapshot

The RVGSnapshot agent depends on these resources.

**Figure 8-3** Sample service group for an RVGSnapshot resource



## Agent functions for RVGSnapshot

The RVGSnapshot agent has the following agent functions:

Online	Creates a transactionally consistent snapshot of all volumes in the RVG.
Offline	Destroys the snapshot.
Monitor	No operation; failure of the snapshot will be indicated by the failure of the Mount resource of any file systems mounted on it.
Clean	Cleans up any failed snapshot creation or deletion.

## State definitions for RVGSnapshot

The RVGSnapshot agent has the following state definitions:

ONLINE	Indicates that a snapshot was created.
OFFLINE	Indicates that a snapshot was destroyed.
FAULTED	The RVGSnapshot resource faults on timeout if a snapshot creation did not succeed during an online.

## Attributes for RVGSnapshot

**Table 8-4** Required attributes

Required attributes	Description
RvgResourceName	The name of the VCS RVG-type resource that manages the RVG that will be snapshot by this agent. Type and dimension: string-scalar
CacheObj	Name of the cache object that is required for a space-optimized snapshot; the fdsetup wizard will create one if it does not exist Type and dimension: string-scalar
Prefix	Token put before the name of the actual volume when creating the snapshotted volumes. Type and dimension: string-scalar

**Table 8-5** Optional attributes

Optional attributes	Description
DestroyOnOffline	A flag to indicate whether to destroy the snapshot upon taking the resources offline. For a fire drill, the snapshot should be deleted to reduce any performance impact of leaving the snapshot for a long period of time; however, if there is interest in keeping the data, then this value should be set to 0. The default is 1 (true). Type and dimension: integer-scalar Default: 1
FDFile	The fire drill schedule updates this attribute with the system name and the path to a file containing the output of the last complete fire drill for the group containing an RVGSnapshot resource. Type and dimension: string-scalar

## Resource type definitions for RVGSnapshot

The resource type definition for the RVGSnapshot agent follows.

```
type RVGSnapshot (
    static keylist RegList = { Prefix }
    static int NumThreads = 1
    static str ArgList[] = { RvgResourceName, CacheObj, Prefix,
        DestroyOnOffline }
```

```
    str RvgResourceName
    str CacheObj
    str Prefix
    boolean DestroyOnOffline = 1
    temp str FDFile
)
```

## Sample configurations for RVGSnapshot

```
RVGSnapshot rvg-sos (
    RvgResourceName = ApplicationRVG
    CacheObj = cacheobj
    Prefix = snap
)
```

## RVGShared agent

Monitors the RVG in a shared environment. This is a parallel resource. The RVGShared agent enables you to configure parallel applications to use an RVG in a cluster. The RVGShared agent monitors the RVG in a shared disk group environment. The RVGShared agent must be configured as a parallel group in VCS. Typically, the RVGShared resource is online or offline at the same time on all the nodes in the VCS cluster. An example configuration file for this agent that can be used as a guide when creating your configuration is located at `/etc/VRTSvcs/conf/sample_vvr/RVGLogowner`.

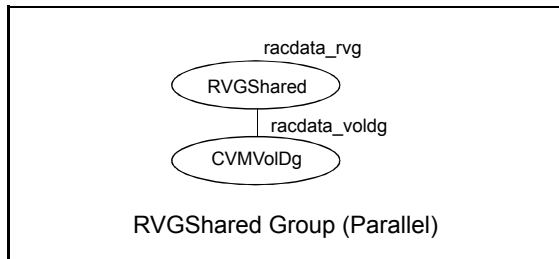
## Dependencies for RVGShared agent

The RVGShared resource represents the RVG of the RDS. The RVGShared resource is dependent on the CVMVolDg resource.

The RVGShared resource must be configured in a parallel group.

Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for information on configuring parallel applications for highly availability.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information on dependencies.

**Figure 8-4** Sample service group for an RVGShared resource

---

**Note:** Do not add any volumes that are part of the RVG in the CVMVolume attribute of the CVMVolDg resource. The volumes in the RVG are managed by the RVGShared resource.

---

## Agent functions for RVGShared agent

The RVGShared agent has the following agent functions:

Online	Verifies whether the RVG is started. If the RVG is not started, recovers and starts the RVG.
Offline	No action.
Monitor	Displays the state as <code>ONLINE</code> if the RVG is started. Displays the state as <code>OFFLINE</code> if the RVG is not started.
Clean	No action.
Info	The info entry point displays information about the replication status of a RDS.

## State definitions for RVGShared agent

The RVGShared agent has the following state definitions:

<code>ONLINE</code>	Indicates that the RVG is in the <code>ENABLED/ACTIVE</code> state.
<code>OFFLINE</code>	Indicates that the RVG is not in the <code>ENABLED/ACTIVE</code> state or that the administrator has invoked the offline entry point.

## Attributes for RVGShared agent

**Table 8-6** Required attributes

Required attributes	Description
RVG	The name of the RVG being monitored. Type and dimension: string-scalar
DiskGroup	The shared-disk group with which this RVG is associated. Type and dimension: string-scalar

## Resource type definitions for RVGShared agent

The RVGShared resource type definition follows.

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

## Sample configurations for RVGShared agent

```
RVGShared racdata_rvg (
    RVG = racl_rvg
    DiskGroup = oradatadg
)
```

## RVGLogowner agent

Assigns and unassigns a node as the logowner in the CVM cluster; this is a failover resource. The RVGLogowner agent assigns or unassigns a node as a logowner in the cluster. To replicate data, VVR requires network connectivity between the Primary and the Secondary. In a shared disk group environment, only one node, that is, the logowner, can replicate data to the Secondary.

For replication to be highly available, the logowner must be highly available. To make the logowner highly available, the RVGLogowner resource must be configured as a resource in a failover group. Also, a virtual IP must be set up on the logowner to enable replication and failover of the logowner from one node to another in a cluster. The virtual IP must be configured as an IP resource.



For more information about the logowner, see the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*. An example configuration file for this agent that can be used as a guide when creating your configuration, is located at `/etc/VRTSvcs/conf/sample_vvr/RVGLogowner`.

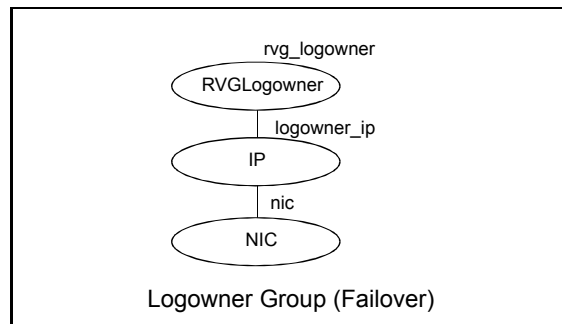
## Dependencies for RVGLogowner agent

The RVGLogowner resource represents the logowner for RVG in the cluster. The RVGLogowner resource is dependent on the IP resource that it uses for replication.

The RVGLogowner resource must be configured in a failover group. The RVGLogowner group is used in conjunction with the RVGSharedPri and RVGShared agents in separate groups, with the appropriate service group dependencies.

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

**Figure 8-5** Sample service group for an RVGLogowner resource



## Agent functions RVGLogowner agent

The RVGLogowner agent has the following agent functions:

- |         |  |
|---------|--|
| Online  | Assigns the logowner on the node.  |
| Offline | Unassigns the logowner on the node.  |
| Monitor | Returns <code>ONLINE</code> if the node is the logowner and the RVG is in <code>ENABLED/ACTIVE</code> state. Returns <code>OFFLINE</code> if the node is the logowner and the state is not <code>ENABLED/ACTIVE</code> , or if the node is not the logowner (regardless of the state). The RVG for which the logowner is monitored must be configured as the <code>RVGShared</code> resource type. |
| Clean   | Unassigns the logowner on the node.  |

## State definitions RVGLogowner agent

The RVGLogowner agent has the following state definitions:

- ONLINE      Indicates that the node is the logowner for the RVG in the cluster.
- OFFLINE     Indicates that the node is not the logowner for the RVG in the cluster.

## Attributes RVGLogowner agent

**Table 8-7**      Required attributes

Required attributes	Description
RVG	The name of the RVG being monitored. Type and dimension: string-scalar Example: "hr_rvg"
DiskGroup	The disk group with which this RVG is associated. Type and dimension: string-scalar Example: "hrbg"

**Table 8-8**      Internal attributes

Bunker attributes	Description
StorageDG	For internal use only, do not modify. The name of the bunker disk group. Type and dimension: string-scalar Example: "hr_bdg"
StorageRVG	For internal use only, do not modify. The name of the bunker RVG. Type and dimension: string-scalar Example: "hr_brvg"
StorageHostIds	For internal use only, do not modify. A space-separated list of the hostids of each node in the bunker cluster. Type and dimension: string-keylist Example: "bunker_host"

## Resource type definitions RVGLogowner agent

The RVGLogowner resource type definition follows.

```
type RVGLogowner (  
    static int NumThreads = 1  
    static str ArgList[] = { RVG, DiskGroup }  
    static int OnlineRetryLimit = 5  
    str RVG  
    str DiskGroup  
    str StorageRVG  
    str StorageDG  
    str StorageHostIds  
)
```

## RVGLogowner agent notes RVGLogowner agent

The RVGLogowner agent has the following notes:

### **CVM master node needs to assume the logowner role for VCS managed VVR resources**

If you use VCS to manage RVGLogowner resources in an SFCFSHA environment or an SF Oracle RAC environment, Symantec recommends that you perform the following procedures. These procedures ensure that the CVM master node always assumes the logowner role. Not performing these procedures can result in unexpected issues that are due to a CVM slave node that assumes the logowner role.

For a service group that contains an RVGLogowner resource, change the value of its PreOnline trigger to 1 to enable it.

#### **To enable the PreOnline trigger from the command line on a service group that has an RVGLogowner resource**

- ◆ On each node in the cluster, perform the following command:

```
# hagrpl -modify RVGLogowner_resource_sg PreOnline 1 -sys system
```

Where *RVGLogowner\_resource\_sg* is the service group that contains the RVGLogowner resource. The *system* is the name of the node where you want to enable the trigger.

On each node in the cluster, merge the preonline\_vvr trigger into the default triggers directory.

### To merge the preonline\_vvr trigger

- ◆ On each node in the cluster, merge the preonline\_vvr trigger to the /opt/VRTSvcs/bin/triggers directory.

```
# cp /opt/VRTSvcs/bin/sample_triggers/VRTSvcs/preonline_vvr \  
/opt/VRTSvcs/bin/triggers
```

Refer to the sample configurations directory for samples of how to enable these triggers (/opt/VRTSvcs/bin/sample\_triggers/VRTSvcs.)

---

**Note:** In case you perform a master switch using the `vxclustadm setmaster` command, the RVGLogowner service group will continue to remain online on the original master, even when the PreOnline attribute is enabled.

Run the following command in order to switch the RVGLogowner service group to the new master node:

```
# hagr -switch RVGLogowner_resource_sg -to new_cvm_master
```

Where `RVGLogowner_resource_sg` is the name of the service group that contains the RVGLogowner resource and `new_cvm_master` is the name of the new CVM master node.

---

## Sample configurations RVGLogowner agent

```
RVGLogowner vvr_rvglogowner (  
  RVG = app_rvg  
  DiskGroup = vvrdbg  
)
```

## RVGSharedPri agent

Attempts to migrate or takeover a Secondary to a Primary when a parallel service group fails over. The RVGSharedPri agent enables migration and takeover of a VVR replicated data set in parallel groups in a VCS environment. Bringing a resource of type RVGSharedPri online causes the RVG on the local host to become a primary if it is not already. The agent is useful when hosts in both the primary and secondary side are clustered using a VCS global cluster, to completely automate the availability of writable replicated disks to an application managed by VCS.

You cannot use the primary-elect feature with this agent. For a detailed description of the primary-elect feature, see *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

The RVGSharedPri agent includes the following key features:

- Removes manual steps of migrating a VVR primary and secondary roles when failing over applications across a wide area.
- Minimizes the need for resynchronizing replicated volumes by attempting a migration before attempting a hard takeover.
- Waits for the two sides of a replicated data set to become completely synchronized before migrating roles.
- Supports an automatic fast fallback resynchronization of a downed primary if it later returns after a takeover.

Sample configuration files are located in the `/etc/VRTSvcs/conf/sample_rac/` directory and include CVR in the filename. These sample files are installed as part of the VRTSdbac RPM, and can be used as a guide when creating your configuration.

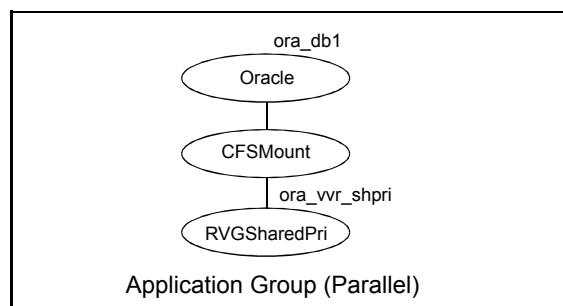
## Dependencies for RVGSharedPri agent

The RVGSharedPri agent is used in conjunction with the RVGShared and RVGLogowner agents in separate groups, with the appropriate service group dependencies.

Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for information on configuring parallel applications for highly availability.

The RVGSharedPri agent must be configured in a parallel service group. The application service group contains the resources managing the actual application and file systems as well as the RVGSharedPri agent.

**Figure 8-6** Sample service group for an RVGSharedPri resource



## Agent functions for RVGSharedPri agent

The RVGSharedPri agent has the following agent functions:

Online	Determines the current role of the RVG; if Secondary, attempt a migrate, waiting for any outstanding writes from the original Primary; if the original Primary is down attempt a takeover; if the RVG is a Primary, perform no actions and go online
Offline	Performs no actions.
Monitor	Performs no actions; monitoring of the actual RVG is done by the RVGShared agent.
Clean	Performs no actions.
fbsync	This is an action entry point.  It resynchronizes the original Primary with the new Primary that has taken over with fast-failback, after the original Primary had become unavailable.  This needs to be executed when the original Primary becomes available and starts acting as a Secondary.
resync	This is an action entry point.  It resynchronizes the Secondaries with the Primary using DCM.

## State definitions for RVGSharedPri agent

The RVGSharedPri agent has the following state definitions:

FAULTED	Monitoring of the actual RVG is done by the RVGShared agent; accidental migration of a VVR Primary outside of VCS would cause other resources to fault immediately, such as Mount, so no special monitoring by this agent is necessary.
---------	---

## Attributes for RVGSharedPri agent

**Table 8-9** Required attributes

Required attributes	Description
RvgResourceName	The name of the RVGShared resource type that this agent will promote, that is, the name RVG resource type which has been configured using the RVGShared agent.  Type and dimension: string-scalar

**Table 8-9** Required attributes (*continued*)

Required attributes	Description
AutoTakeover	A flag to indicate whether the agent should perform a takeover on online if the original Primary is down.  Type and dimension: integer-scalar  Default: 1
AutoResync	A flag to indicate whether the agent should attempt to automatically perform a fast-failback resynchronization of the original Primary after a takeover and after the original Primary returns.  Type and dimension: integer-scalar  Default: 0
VCSResLock	This attribute is reserved for internal use by VCS.  Type and dimension: string-scalar

## Resource type definitions for RVGSharedPri agent

The RVGSharedPri resource type definition follows.

```

type RVGSharedPri (
    static keylist SupportedActions = { fbsync, resync }
    static int NumThreads = 1
    static int OnlineRetryLimit = 1
    static str ArgList[] = { RvgResourceName, "RvgResourceName:RVG",
        "RvgResourceName:DiskGroup", AutoTakeover, AutoResync }
    str RvgResourceName
    int AutoTakeover = 1
    int AutoResync = 0
    temp str VCSResLock
)

```

## Sample configurations for RVGSharedPri agent

```

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

```

# Index

## Symbols

802.1Q trunking 79

## A

about

Network agents 78

Samba agents 145

agent

modifying 18

agent functions 142

Apache HTTP server agent 160

Application agent 175

CoordPoint agent 185

DiskGroup agent 23

DiskGroupSnap agent 34

DiskReservation agent 47

DNS agent 114

ElifNone agent 225

FileNone agent 227

FileOnOff agent 229

FileOnOnly agent 231

IP agent 80

IPMultiNIC agent 93

LVMLogicalVolume agent 57

LVMVolumeGroup agent 60

Mount agent 66

MultiNICA agent 101

Linux 101

NetBIOS agent 155

NFS agent 131

Linux 131

NFSRestart agent 136

Linux 136

NIC agent 87

NotifierMngr agent 205

Phantom agent 215

Process agent 196

ProcessOnOnly agent 200

Proxy agent 211

RemoteGroup agent 217

SambaServer agent 148

agent functions *(continued)*

SambaShare agent 152

Share agent 142

Volume agent 52

VolumeSet agent 54

agents. *See* RVG agent

Apache HTTP server 159

Application 173

CoordPoint 184

DiskGroup 22

DiskGroupSnap 32

DiskReservation 47

DNS 113

ElifNone 224

FileNone 226

FileOnOff 228

FileOnOnly 230

IP 79

IPMultiNIC 92

LVMLogicalVolume 56

LVMVolumeGroup 60

Mount 65

MultiNICA 99

Linux 99

NetBIOS 154

NFS 130

NFSRestart 135

NIC 85

NotifierMngr 204

Phantom 214

Process 195

ProcessOnOnly 200

Proxy 211

RemoteGroup 216

RVGLogowner. *See* RVGLogowner agent

RVGPrimary. *See* RVGPrimary agent

RVGShared. *See* RVGShared agent

RVGSharedPri. *See* RVGSharedPri agent

RVGSnapshot. *See* RVGSnapshot agent

SambaServer 148

SambaShare 152

Share 141



**agents** (*continued*)

Volume 51

**Apache HTTP server agent**

agent functions 160

attributes 162

description 159

detecting application failure 167

sample configuration 169

state definitions 161

**Application agent**

agent functions 175

attributes 178

description 173

high availability fire drill 174

Linux attributes 178

resource type definition 181

sample configurations 183

state definitions 177

**association dimension** 18**attribute data types** 18**attributes**

Application agent 178

Linux 178

CoordPoint agent 186

DiskGroup agent 25

Linux 25

DiskGroupSnap agent 34

DiskReservation agent 48

Linux 48

DNS agent 116

ElifNone agent 225

FileNone agent 228

FileOnOff agent 230

FileOnOnly agent 232

IP agent 81

Linux 81

IPMultiNIC agent 94

Linux 94

LVMLogicalVolume agent 58

Linux 58

LVMVolumeGroup agent 61

Linux 61

modifying 18

Mount agent 69

Linux 69

MultiNICA agent 102

Linux 102

NFS agent 132

Linux 132

**attributes** (*continued*)

NFSRestart agent 138

Linux 138

NIC agent 88

Linux 88

NotifierMngr agent 206

Linux 206

Process agent 198

Linux 198

ProcessOnOnly 201

ProcessOnOnly agent

Linux 201

Proxy agent 212

RemoteGroup agent 218

SambaServer agent 149

Share agent 143

Linux 143

Volume agent 53

VolumeSet agent 55

**AutoResync attribute**

RVGPrimary agent 255

**B**

bonded network interfaces 86

boolean data types 18

**C**

Cluster Manager (Java Console)

modifying attributes 18

CNAME record 125

configuration files

main.cf 215

modifying 18

types.cf 18

CoordPoint agent

agent functions 185

attributes 186

description 184

resource type definition 186

sample configurations 187

state definitions 185

**D**

data type

boolean 18

string 18

data types

integer 18

- dependency graphs
    - RVGLogowner agent 249
    - RVGPrimary agent 237
    - RVGShared agent 246
    - RVGSharedPri agent 253
  - description
    - resources 18
  - dimensions
    - keylist 18
    - scalar 18
    - vector 18
  - DiskGroup agent
    - agent functions 23
    - attributes 25
    - description 22
    - high availability fire drill 30
    - Linux attributes 25
    - resource type definition 29
    - sample configurations 32
    - state definitions 25
  - DiskGroupSnap agent
    - agent functions 34
    - attributes 34
    - description 32
    - resource type definition 40
    - sample configurations 40
    - state definitions 34
  - DiskReservation agent
    - agent functions 47
    - attributes 48
    - description 47
    - Linux attributes 48
    - resource type definition 49
    - sample configurations 50
    - state definitions 47
  - DNS agent 115
    - agent functions 114
    - attributes 116
    - description 113
    - resource type definition 122
    - sample web server configuration 125
- E**
- ElifNone agent
    - agent functions 225
    - attributes 225
    - description 224
    - resource type definition 226
    - sample configuration 226
  - ElifNone agent (*continued*)
    - state definitions 225
- F**
- failover group
    - RVGLogowner agent 248
  - fast failback
    - AutoResync attribute of RVGPrimary 255
  - fast failback resynchronization
    - RVGPrimary 237
    - RVGSharedPri 253
  - FileNone agent
    - agent functions 227
    - attribute 228
    - description 226
    - resource type definition 228
    - sample configurations 228
    - state definitions 227
  - FileOnOff agent
    - agent functions 229
    - attribute 230
    - description 228
    - state definitions 229
  - FileOnOnly agent
    - agent functions 231
    - attribute 232
    - description 230
    - resource type definition 232
    - sample configuration 232
    - state definitions 231
  - fire drill
    - RVGSnaphot agent 243
- H**
- high availability fire drill 30, 75, 80, 86, 125, 139, 174, 196
- I**
- integer data types 18
  - IP agent
    - agent functions 80
    - attributes 81
    - description 79
    - high availability fire drill 80
    - Linux attributes 81
    - resource type definitions 84
    - sample configurations 85
    - state definitions 81

IPMultiNIC agent  
 agent functions 93  
 attributes 94  
 description 92  
 Linux attributes 94  
 resource type definitions 95  
 sample configuration 96  
 state definitions 93

## K

keylist dimension 18

## L

logowner  
 virtual IP requirement 248

LVMLogicalVolume agent  
 agent functions 57  
 attributes 58  
 description 56  
 Linux attributes 58  
 resource type definition 58  
 sample configurations 58  
 state definitions 57

LVMVolumeGroup agent  
 agent functions 60  
 attributes 61  
 description 60  
 Linux attributes 61  
 resource type definition 62  
 state definitions 61

## M

main.cf 18, 215  
 main.xml 18  
 migrating  
 RVGPrimary 237  
 RVGSharedPri 252

modifying  
 configuration files 18

modifying agents 18

monitor scenarios  
 DNS agent 125

monitoring bonded NICs  
 Linux 90

Mount agent  
 agent functions 66, 68  
 attributes 69  
 description 65

Mount agent (*continued*)  
 high availability fire drill 75, 125, 139  
 Linux attributes 69  
 notes 74  
 resource type definition 73  
 sample configurations 77

MultiNICA agent  
 agent functions 101  
 Linux 101  
 attributes 102  
 description 99  
 IP Conservation mode 100  
 Linux  
 attributes 102  
 description 99  
 IP Conservation mode 100  
 Performance Mode 101  
 Performance mode 101  
 resource type definitions 107  
 resource type definitions, Linux 107

## N

NetBIOS agent  
 agent functions 155  
 description 154  
 resource type definition 156  
 sample configurations 157  
 state definitions 155

NFS agent  
 agent functions 131  
 Linux 131  
 attributes 132  
 Linux 132  
 description 130  
 resource type definition 133  
 sample configurations 134  
 state definitions 132

NFS lock recovery 133

NFSRestart agent  
 agent functions 136  
 Linux 136  
 attributes 138  
 Linux 138  
 description 135  
 resource type definition 138  
 sample configuration 140  
 state definitions 137

NIC agent  
 agent functions 87

**NIC agent** (*continued*)

- attributes 88
  - Linux 88
- description 85
- high availability fire drill 86
- resource type definitions
  - Linux 90
- sample configurations 91
- state definitions 87

**noautoimport flag** 31**NotifierMngr agent**

- agent functions 205
- attributes 206
- description 204
- Linux attributes 206
- resource type definition 208
- sample configurations 209
- state definitions 205

**O**

online query 125

**P****parallel group**

- RVGShared agent 246

**Phantom agent**

- agent functions 215
- description 214
- resource type definition 215
- sample configurations 215

**prerequisites**

- NFS lock recovery 133
  - Linux 133
- Samba agents 146

**Process agent**

- agent functions 196
- attributes 198
- description 195
- high availability fire drill 196
- Linux attributes 198
- resource type definition 199
- sample configurations 199
- state definitions 197

**ProcessOnOnly agent**

- agent functions 200
- attributes 201
- description 200
- Linux attributes 201

**ProcessOnOnly agent** (*continued*)

- resource type definition 202
- sample configurations 203
- state definitions 201

**Proxy agent**

- agent functions 211
- attributes 212
- description 211
- resource type definition 212
- sample configurations 212

**R****RemoteGroup agent**

- agent functions 217
- attributes 218
- description 216
- resource type definition 223
- state definitions 217

**resource type definition** 53

SambaShare agent 153

**resource type definitions**

- Application agent 181
- CoordPoint agent 186
- DiskGroup agent 29
- DiskGroupSnap agent 40
- DiskReservation agent 49
- DNS agent 122
- ElifNone agent 226
- FileNone agent 228
- FileOnOnly agent 232
- IP agent 84
- IPMultiNIC agent 95
- LVMLogicalVolume agent 58
- LVMVolumeGroup agent 62
- Mount agent 73
- MultiNICA agent 107
- MultiNICA agent, Linux 107
- NetBIOS agent 156
- NFS agent 133
- NFSRestart agent 138
- NIC agent 90
- NotifierMngr agent 208
- Phantom agent 215
- Process agent 199
- ProcessOnOnly agent 202
- Proxy agent 212
- RemoteGroup agent 223
- SambaServer agent 151
- Share agent 144

- resource type definitions *(continued)*
  - Volume agent 53
- resources
  - description of 18
- RVG agent
  - described 233
- RVGLogowner agent
  - dependency graph 249
  - described 248
  - failover group 248
- RVGPrimary agent
  - dependency graph 237
  - described 237
  - migrating 237
  - takeover 237
- RVGShared agent
  - dependency graph 246
  - described 246
  - parallel group 246
- RVGSharedPri agent
  - dependency graph 253
  - described 252
  - migrating 252
  - takeover 252
- RVGSnapshot agent
  - described 243
  - fire drill 243

## S

- Samba agents 146
  - overview 145
  - prerequisites 146
- SambaServer agent
  - agent functions 148
  - attributes 149
  - description 148
  - resource type definition 151
  - sample configuration 151
  - state definitions 149
- SambaShare agent 152
  - agent functions 152
  - attributes 153
  - resource type definition 153
  - sample configurations 153
  - state definitions 152
- sample configurations
  - Apache HTTP server agent 169
  - Application agent 183
  - CoordPoint agent 187

- sample configurations *(continued)*
  - DiskGroup agent 32
  - DiskGroupSnap agent 40
  - DiskReservation agent 50
  - ElifNone agent 226
  - FileNone agent 228
  - FileOnOff agent 230
  - FileOnOnly agent 232
  - IP agent 85
  - IPMultiNIC 96
  - LVMLogicalVolume agent 58
  - Mount agent 77
  - NetBIOS agent 157
  - NFS agent 134
  - NFSRestart agent 140
  - NIC agent 91
  - NotifierMngr agent 209
  - Phantom agent 215
  - Process agent 199
  - ProcessOnOnly agent 203
  - Proxy agent 212
  - SambaServer agent 151
  - SambaShare agent 153
  - Share agent 145
  - Volume agent 53
- scalar dimension 18
- secure DNS update 126
- setting Mii and miimon 90
- Share agent
  - agent functions 142
    - Linux 142
  - attributes 143
  - description 141
  - Linux 142
  - Linux attributes 143
  - resource type definitions 144
  - sample configurations 145
  - state definitions 142
- snapshots
  - using RVGSnapshot agent for 243
- State definitions
  - VolumeSet agent 55
- state definitions 115
  - Apache HTTP server agent 161
  - Application agent 177
  - CoordPoint agent 185
  - DiskGroup agent 25
  - DiskGroupSnap agent 34
  - DiskReservation agent 47

state definitions (*continued*)

- DNS agent 115
- ElifNone agent 225
- FileNone agent 227
- FileOnOff agent 229
- FileOnOnly agent 231
- IP agent 81
- IPMultiNIC agent 93
- LVMLogicalVolume agent 57
- LVMVolumeGroup agent 61
- Mount agent 68
- NetBIOS agent 155
- NFS agent 132
- NFSRestart agent 137
- NIC agent 87
- NotifierMngr agent 205
- Process agent 197
- ProcessOnOnly agent 201
- RemoteGroup agent 217
- SambaServer agent 149
- SambaShare agent 152
- Share agent 142
- Volume agent 52

string data type 18

**T**

## takeover

- RVGPrimary 237
- RVGSharedPri 252

trunking 79

types.cf 18

**V**

## VCS

- resource types 18

vector dimension 18

## virtual IP

- RVGLogowner agent requirement 248

## Volume agent

- agent functions 52
- attributes 53
- description 51
- sample configurations 53
- state definitions 52

volume sets 31

## VolumeSet agent

- agent functions 54
- attributes 55

VolumeSet agent (*continued*)

State definitions 55