

Veritas[™] Cluster Server One Agent Developer's Guide

AIX, ESX, HP-UX, Linux, Solaris, Windows

5.0 Service Pack 1



Veritas Cluster Server One Agent Developer's Guide

Copyright © 2010 Symantec Corporation. All rights reserved.

Product version: Veritas Cluster Server One5.0 Service Pack 1
Document version: 5.0.SP1.0

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

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

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

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

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

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

Technical Support

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

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

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

www.symantec.com/business/support

Contacting Technical Support

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

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

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support

Customer Service is available to assist with the following types of issues:

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

Documentation feedback

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

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

Maintenance agreement resources

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

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

Additional enterprise services

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

Enterprise services that are available include the following:

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

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

www.symantec.com/business/services/

Select your country or language from the site index

Contents

Chapter 1	Introduction	
	About VCS One agents	14
	How agents work	15
	About the agent framework	15
	Resource type definitions	15
	About agent functions (entry points)	16
	About on-off, on-only, and persistent resources	16
	About attributes	17
	About intentional offline of applications	21
	About developing an agent	22
	Considerations for the application	22
	High-level overview of the agent development process	23
Chapter 2	Agent entry point overview	
	About agent entry points	26
	Supported entry points	26
	How the agent framework interacts with entry points	26
	Agent entry points described	27
	About the monitor entry point	27
	About the info entry point	28
	About the online entry point	30
	About the offline entry point	30
	About the clean entry point	31
	About the action entry point	33
	About the attr_changed entry point	34
	About the open entry point	34
	About the close entry point	35
	About the shutdown entry point	35
	Return values for entry points	36
	Considerations for using C++ or script entry points	37
	About the VCSAgStartup routine	37
	About the ArgList and ArgListValues attributes	39
	ArgListValues attribute for agents registered as V50 and later	39
	Overview of the name-value tuple format	39
	About the entry point timeouts	40

ArgListValues attribute for agents registered as V40 and earlier 41

Chapter 3 Creating entry points in C++

About creating entry points in C++	46
Entry point examples in this chapter	47
Data Structures	48
Syntax for C++ entry points	49
Syntax for C++ VCSAgStartup	49
Syntax for C++ monitor	50
Syntax for C++ info	51
Syntax for C++ online	55
Syntax for C++ offline	56
Syntax for C++ clean	57
Syntax for C++ action	58
Syntax for C++ attr_changed	60
Syntax for C++ open	62
Syntax for C++ close	62
Syntax for C++ shutdown	64
Agent framework primitives	65
VCSAgRegisterEPStruct	65
VCSAgSetCookie	65
VCSAgSetCookie2	65
VCSAgRegister	67
VCSAgUnregister	68
VCSAgGetCookie	69
VCSAgStrlcpy	70
VCSAgStrlcat	70
VCSAgSnprintf	70
VCSAgCloseFile	70
VCSAgDelString	70
VCSAgExec	71
VCSAgExecWithTimeout	72
VCSAgSendTrap	73
VCSAgLockFile	73
VCSAgInitEntryPointStruct	75
VCSAgSetStackSize	75
VCSAgUnlockFile	75
VCSAgDisableCancellation	76
VCSAgRestoreCancellation	76
VCSAgSetEntryPoint	77
VCSAgValidateAndSetEntryPoint	77
VCSAgSetLogCategory	77
VCSAgGetProductName	77

VCSAgIsMonitorLevelOne	77
VCSAgIsMonitorLevelTwo	77
VCSAgMonitorReturn	78
VCSAgSetResEPTimeout	78
VCSAgDecryptKey	78
VCSAgCloseFile	78
VCSAgCloseFilePtr	79
VCSAgCloseFp	79
VCSAgCloseFpPtr	80
VCSAgCloseDirp	80
VCSAgCloseDirpPtr	80
VCSAgDelPasswdPtr	81
VCSAgDelString	81
VCSAgDelStringPtr	81
Agent Framework primitives for container support	82
VCSAgISContainerCapable	82
VCSAgExecInContainerWithTmo	82
VCSAgExecInContainerWithTimeout	82
VCSAgGetUID	83
VCSAgIsPidInContainer	83
VCSAgIsProcInContainer	83
VCSAgGetContainerID2	84
VCSAgGetContainerName2	84
Agent framework primitives for Windows	85
C++ logging primitives for use by Windows agents	85
C++ primitives for ArgList conversion	86
C++ primitives for ArgList string manipulation	87
C++ primitives for allocating and freeing memory	88
Example: Windows-based primitives for the FileOnOff agent	88

Chapter 4

Creating entry points in scripts

About creating entry points in scripts	94
Rules for using script entry points	94
Parameters and values for script entry points	95
ArgList attributes	95
Example	95
Syntax for script entry points	97
Syntax for the monitor script	97
Syntax for the online script	97
Syntax for the offline script	97
Syntax for the clean script	98
Syntax for the action script	98
Syntax for the attr_changed script	98

Syntax for the info script	98
Syntax for the open script	99
Syntax for the close script	99
Syntax for the shutdown script	99
Example script entry points	100

Chapter 5 Logging agent messages

About logging agent messages	106
Logging in C++ and script-based entry points	106
Agent messages: format	107
C++ agent logging APIs	108
Agent application logging macros for C++ entry points	109
Agent debug logging macros for C++ entry points	110
Severity arguments for C++ macros	111
Initializing function_name using VCSAG_LOG_INIT	112
Log category	113
Examples of logging APIs used in a C++ agent	114
Script entry point logging functions	117
Using functions in scripts	118
VCSAG_SET_ENVS	118
VCSAG_LOG_MSG	121
VCSAG_LOGDBG_MSG	122
Example of logging functions used in a script agent	124

Chapter 6 Building a custom agent

Files for use in agent development	126
Creating the type definition file for a custom agent	127
Naming convention for the type definition file	127
Example: Type definition file for the FileOnOff agent on Linux	127
Example: Type definition file for the FileOnOff agent on Windows	128
Requirements for creating the agentTypes.xml file	129
Configuring attribute properties	129
Example: Type definition for a custom agent that supports intentional offline	129
Adding the custom type definition to the configuration	130
Building a custom agent on UNIX	131
Example: Using script entry points on UNIX	132
Example: Using VCSAgStartup() and script entry points on UNIX	134
Implementing entry points using C++	135
Example: Using C++ entry points on UNIX	135
Example: Using C++ and script entry points on UNIX	138
..... Building a custom agent on Windows	141

Recommendations for building custom agents on Windows	141
Example: Using C++ entry points on Windows	142
.....Implementing entry points using scripts	145
Example: Using script entry points on Windows	145
Example: Using VCSAgStartup() and script entry points on Windows	146
Example: Using C++ and script entry points on Windows	147
Installing the custom agent	150
Defining resources for the custom resource type	150
Sample resource definition	150

Chapter 7 Testing agents

About testing agents	154
Using debug messages	154
Debugging agent functions (entry points).	154
Debugging the agent framework	155
Using the engine process to test agents	155
Test commands	156
Using the AgentServer utility to test agents	157

Chapter 8 Static type attributes

About static attributes	162
Overriding static type attributes	162
Static timeout attributes on Windows	163
Static type attribute definitions	164
ActionTimeout	164
AEPTIMEOUT	164
AgentClass	164
AgentFailedOn	164
AgentPriority	165
AgentReplyTimeout	165
AgentStartTimeout	165
ArgList	165
AttrChangedTimeout	166
CleanTimeout	166
CleanRetryLimit	166
CloseTimeout	167
ContainerOpts	167
ConfInterval	168
FaultOnMonitorTimeouts	169
InfoInterval	169
InfoTimeout	170
IntentionalOffline	170

- Leveltwomonitorfrequency 170
- LogDbg 170
- LogFileSize 171
- MinCompatAgentVersion 171
- MonitorInterval 172
- MonitorMethod 172
- MonitorStatsParam 172
- MonitorTimeout 172
- NumThreads 173
- OfflineMonitorInterval 173
- OfflineTimeout 173
- OnlineRetryLimit 174
- OnlineTimeout 174
- OnlineWaitLimit 174
- OpenTimeout 174
- Operations 175
- Platform 175
- RegList 175
- ResFaultPolicy 177
- RestartLimit 177
- ScriptClass 178
- ScriptPriority 178
- SupportedActions 179
- ToleranceLimit 180
- TypeDefinitionVersion 180

Chapter 9 State transition diagrams

- State transitions 182
- State transitions for quiesce and unquiesce 191
- State transitions with respect to ResFaultPolicy attribute 201

Appendix A Using VCS agents with VCS One

- About using VCS agents with VCS One 210
- Using pre-5.0 VCS agents with VCS One and registering them as V51 210
 - Outline of steps to change V40 agents V51 210
 - Sourcing ag_i18n_inc modules in script entry points 211

- Index 213

Introduction

- [About VCS One agents](#)
- [How agents work](#)
- [About developing an agent](#)

About VCS One agents

Agents are programs that manage resources, such as a disk group or a mount point, within a cluster environment. Each type of resource requires an agent. The agent acts as an intermediary between VCS One and the resources it manages, typically by bringing it online, monitoring its state, or taking it offline.

Agents packaged with VCS One are referred to as *bundled agents*. Examples of bundled agents include the IP (Internet Protocol) and NIC (network interface card) agents. For more information on VCS One bundled agents, including their attributes and modes of operation, see the *Bundled Agents Reference Guide*.

Agents packaged separately from the product for use with VCS One are referred to as *high availability agents*. They include agents for Oracle, for example. Contact your Symantec sales representative for information on how to purchase these agents for your configuration.

Note: Custom agents, that is, agents developed outside of Symantec, are not supported by Symantec Technical Support.

How agents work

A single agent can manage multiple resources of the same type on one system. For example, the NIC agent manages all NIC resources. The resources to be managed are those defined within the VCS One configuration.

The VCS One daemon that runs on the systems hosting the applications is `vcsonclientd`.

When the VCS One process `vcsonclientd` comes up on a system, it communicates with the Policy Master. The Policy Master instructs `vcsonclientd` to start the configured agents.

An agent carries out the commands from VCS One to bring resources online, monitor their status, and take them offline, as needed. When an agent crashes or hangs, VCS One detects the fault and restarts the agent.

About the agent framework

The agent framework is a multi-threaded dynamically loadable library that contains most of the logic required for an agent to interact with VCS One. The library also exposes several APIs for agent developers to use. It is the framework that is responsible for understanding certain common configuration attributes (such as `MonitorInterval`, `RestartLimit`, and `ToleranceLimit`), invoking agent entry points at the appropriate times, heartbeating with `vcsonclientd` and also receiving commands sent by the VCS One engine (via `vcsonclientd`). When an agent's entry points are written in C++, the compiled agent links dynamically with the agent framework library. When an agent is developed using scripts (script entry points) such as Perl or shell, the provided agent binary (`Script51Agent`) runs as the agent binary and loads the agent framework library when it starts. The agent framework handles much of the complexity that need not concern the agent developer.

Resource type definitions

The agent for each type of resource requires a resource type definition that describes the information an agent needs to control resources of that type. The type definition can be considered similar to a header file in a C program. The type definition defines the data types of attributes to provide error checking, and to provide default values for certain attributes that affect all resources of that resource type.

One of the attributes that is defined for the IP resource type is the `Address` attribute, which stores the actual IP address of a specific IP resource. This attribute is defined as a 'string-scalar' as the use of periods to denote an IP address makes the acceptable values for the attribute a string.

About agent functions (entry points)

An entry point is either a C++ function or a script (shell or Perl, for example) used by the agent to carry out a specific task on a resource. The agent framework supports a specific set of entry points, each of which is expected to do a different task and return. For example, the online entry point brings a resource online.

See “[Supported entry points](#)” on page 26.

An agent developer implements the entry points for a resource type that the agent uses to carry out the required tasks on the resources of that type. For example, in the online entry point for the Mount resource type, the agent developer includes the logic to mount a file system based on the parameters provided to the entry point. These parameters are attributes for a particular resource, for example, mount point, device name, and mount options. In the monitor entry point, the agent developer checks the state of the mount resource and returns a code to indicate whether the mount resource is online or offline.

See “[About agent entry points](#)” on page 26.

About on-off, on-only, and persistent resources

Different types of resources require different types of control, requiring implementation of different entry points. Resources can be classified as on-off, on-only, or persistent.

■ On-off resources

Most resources are on-off, meaning agents start and stop them as required. For example, VCS One assigns an IP address to a specified NIC when bringing a resource online and removes the assigned IP address when taking the resource offline. Another example is the DiskGroup resource. VCS One imports a disk group when needed and deports it when it is no longer needed. For agents of on-off resources, all entry points can be implemented.

■ On-only resources

An on-only resource is brought online, but it is not taken offline when the associated service group is taken offline. For example, in the case of the FileOnOnly resource, the engine creates the specified file when required, but does not delete the file if the associated service group is taken offline. For agents of on-only resources, the offline entry point is not needed or invoked.

■ Persistent resources

A persistent resource cannot be brought online or taken offline, yet the resource must be present in the configuration to allow the resource to be monitored. For example, a NIC resource cannot be started or stopped, but it

is required to be operational in order for the associated IP address to function properly. The agent monitors persistent resources to ensure their status and operation. An agent for a persistent resource does not require or invoke the online or offline entry points. It uses only the monitor entry points.

About attributes

VCS One has the following types of attributes, depending on the object the attribute applies to.

Resource type attributes	Attributes associated with resource types in VCS One. These can be further classified as: ■ Type-independent—Attributes that all agents (or resource types) understand. Examples: RestartLimit and MonitorInterval; these can be set for any resource type. Typically, these attributes are set for all resources of a specific type. For example, if you set the MonitorInterval for the IP resource type, the same value applies to all resources of type IP. You can also override the values of these attributes, that is, you can configure a different attribute value for each resource of this type. ■ Type-dependent—Attributes that apply to a particular resource type. Examples: The MountPath attribute applies only to the Mount resource type. The Address attribute applies only to the IP resource type. Attributes defined in the types file (types.<i>platform</i>.xml) apply to all resources of the resource type. When you configure resources, you can assign resource-specific values to these attributes, which appear in the main.xml file. For example, the PathName attribute for the FileOnOff resource type is type-dependent, and can take a resource-specific value when configured. ■ Static—These attributes apply to all resource types and can have a different value per resource type. You can override some static attributes and assign them resource-specific values. These attributes are prefixed with the term <code>static</code> and are not included in the resource's argument list.
--------------------------	--

Attribute data types

VCS One supports the following data types for attributes.

String	A string is a sequence of characters. If the string contains double quotes, the quotes must be immediately preceded by a backslash. A backslash is represented in a string as \\. Quotes are not required if a string begins with a letter, and contains only letters, numbers, dashes (-), and underscores (_). For example, a string defining a network interface such as hme0 or eth0 does not require quotes as it contains only letters and numbers. However a string defining an IP address contains periods and requires quotes, such as: "192.168.100.1" VCS One also supports UTF-8 encoded values for some attributes.
Integer	Signed integer constants are a sequence of digits from 0 to 9. They may be preceded by a dash, and are interpreted in base 10. Integers cannot exceed the value of a 32-bit signed integer: 21471183247.
Boolean	A boolean is an integer, the possible values of which are 0 (false) and 1 (true).

Attribute dimensions

VCS One attributes have the following dimensions.

Scalar	A scalar has only one value. For example: <pre><attribute name="myattribute" type="str" dimension="scalar"> </attribute></pre> When values are assigned to a scalar attribute in the main.xml configuration file, it resembles: <pre><attribute name="myattribute"> <scalar>5</scalar> </attribute></pre> OR: <pre><attribute name="myattribute"> <scalar>"myattribute_val1"</scalar> </attribute></pre>
--------	--

Vector

A vector is an ordered list of values. Each value is indexed using a positive integer beginning with zero.

Example snippet from the type definition file for an agent:

```
<attribute name="MyVector" type="str"
dimension="vector">
  <default>
    <val>"value1"</val>
    <val>"Value2"</val>
  </default>
</attribute>
```

When values are assigned to a vector attribute in the main.xml configuration file, the attribute definition might resemble:

```
<attribute name="MyVector"
  <val>"new_value1"</val>
  <val>"new_value2"</val>
</attribute>
```

Keylist

A keylist is an unordered list of strings, and each string is unique within the list.

For example:

```
<attribute name="MyKeylist" type="str"
dimension="keylist">
  <default>
    <val key="key1"></val>
  </default>
</attribute>
```

When values are assigned to a keylist attribute in the main.xml file, it might resemble:

```
<attribute name="MyKeylist">
  <val key="KeyName1"></val>
  <val key="KeyName2"></val>
</attribute>
```

Association An association is an unordered list of name-value pairs.

For example:

```
<attribute name="MYAssoc" type="str"
dimension="assoc">
  <default>
    <val key="key1">"value1"</val>
    <val key="key2">"value1"</val>
  </default>
</attribute>
```

When values are assigned to an association attribute in the `main.xml` file, it might resemble:

```
<attribute name="MyAssoc">
  <val key="newKey1">"Value1"</val>
  <val key="newKey2">"Value2"</val>
</attribute>
```

Attribute scope across systems: global and local attributes

An attribute whose value is the same across all systems on which the service group is configured *global* in scope. An attribute whose value applies on a per-system basis is *local* in scope.

You can set the scope for an attribute on per-system using the following expression in the `main.xml` file:

`context = "systemname".`

In the following example of the `FileOnOff` resource definition in the `main.xml` file, the `PathName` attribute applies locally and is defined specifically for each system where it may be online. Localized attributes are indicated by the expression: `context = "systemname"`:

```
<resources>
  <resource name="G1_fileonoff_res" type="FileOnOff">
    <attribute name="PathName" context = "sysa">
      <scalar>"/tmp/cfmdtype13a/G1_fileonoff_res"
    </scalar>
    </attribute>
    <attribute name="PathName" context = "sysb">
      <scalar>"/tmp/cfmdtype14b/G1_fileonoff_res_2"
    </scalar>
    </attribute>
  </resource>
```

A global attribute is defined only once in a resource definition and does not specify a context.

Attribute life: temporary attributes

You can define temporary attributes in the types file. The values of temporary attributes remain in memory as long as the VCS One `vcsooneclientd` process is running. Values of temporary attributes are not available when the `vcsooneclientd` process is restarted.

These attribute values are not stored in the Policy Master database.

Temporary attributes cannot be converted to permanent, and vice-versa. When you save a configuration, VCS One saves the temporary attribute definitions and their default values in the type definition file.

You can modify attribute values only while VCS One is running.

About intentional offline of applications

Certain agents can identify when an application has been intentionally shut down outside of VCS One control.

If an administrator intentionally shuts down an application outside of VCS One control, VCS One does not treat it as a fault. VCS One sets the service group state as offline or partial, depending on the state of other resources in the service group.

This feature allows administrators to stop applications without causing failovers.

See “[IntentionalOffline](#)” on page 170.

About developing an agent

Before creating the agent, some considerations and planning are required, especially regarding the type of the resource for which the agent is being created.

Considerations for the application

The application for which an agent for VCS One is developed must lend itself to being controlled by the agent and be able to operate in a cluster environment. The following criteria describe an application that can successfully operate in a cluster:

- The application must be capable of being started by a defined procedure. There must be some means of starting the application's external resources such as file systems that store databases, or IP addresses used for listener processes, and so on.
- Each instance of an application must be capable of being stopped by a defined procedure. Other instances of the application must not be affected.
- The application must be capable of being stopped cleanly, by forcible means if necessary.
- Each instance of an application must be capable of being monitored uniquely. Monitoring can be simple or in-depth so as to achieve a high level of confidence in the operation of the application. Monitoring an application becomes more effective when the monitoring procedure resembles the actual activity of the application's user.
- For failover capability, the application must be capable of storing data on shared disks rather than locally or in memory, and each system must be capable of accessing the data and all information required to run the application.
- The application must be crash-tolerant. It must be capable of being run on a system that crashes and of being started on a failover node in a known state. This typically means that data is regularly written to shared storage rather than stored in memory.
- The application must be host-independent within a cluster; that is, there are no licensing requirements or host name dependencies that prevent successful failover.
- The application must run properly with other applications in the cluster.

High-level overview of the agent development process

The steps to create and implement an agent are described by example in later chapters.

Developing the entry points

Decide whether to implement the agent entry points using C++ code, scripts, or a combination of the two.

See [“Considerations for using C++ or script entry points”](#) on page 37.

Create the entry points.

See [Chapter 3, “Creating entry points in C++”](#).

See [Chapter 4, “Creating entry points in scripts”](#).

Building the agent

Build the agent, create required files, and place the agent in specific directories.

The types file contains definitions of resource types that come bundled with VCS One. The file name is:

`types.platform.xml`

A custom resource type definition should be placed in a file that specifies the name of the custom resource. for example:

`MyResourceTypesaix.xml`

This file is referenced as an “include” tag in the VCS One configuration file, `main.xml`.

See [“Creating the type definition file for a custom agent”](#) on page 127.

For more information about the resource type definition, see

[Chapter 6, “Building a custom agent”](#) on page 125.

For building an agent, sample files are provided.

Testing the agent

Test the agent using the Agent Server utility or by defining the resource type in a configuration.

See [Chapter 7, “Testing agents”](#) on page 153.

Agent entry point overview

- [About agent entry points](#)
- [Agent entry points described](#)
- [Return values for entry points](#)
- [Considerations for using C++ or script entry points](#)
- [About the ArgList and ArgListValues attributes](#)

About agent entry points

Developing an agent involves developing the *entry points* that the agent can call to perform operations on a resource, such as to bring a resource online, to take a resource offline, or to monitor the resource.

Supported entry points

The agent framework supports the following entry points:

- `monitor` - determines the status of a resource
- `info` - provides information about an online resource
- `online` - brings a resource online
- `offline` - takes a resource offline
- `clean` - terminates ongoing tasks associated with a resource
- `action` - starts a defined action for a resource
- `attr_changed` - responds to a resource's attribute's value change
- `open` - initializes a resource before the agent starts to manage it
- `close` - deinitializes a resource before the agent stops managing it
- `shutdown` - called when the agent shuts down

See [“Agent entry points described”](#) on page 27.

How the agent framework interacts with entry points

The agent framework ensures that only one entry point is running for a given resource at one time. If multiple requests are received or multiple events are scheduled for the same resource, the agent queues them and processes them one at a time. An exception to this behavior is an optimization such that the agent framework discards internally generated *periodic monitoring* requests for a resource that is already being monitored or that has a monitor request pending.

The agent framework is multithreaded. This means a single agent process can run entry points for multiple resources simultaneously. However, if an agent receives a request to take a given resource offline and simultaneously receives a request to close it, it calls the `offline` entry point first. The `close` entry point is called only after the `offline` request returns or times out. If the `offline` request is received for one resource, and the `close` request is received for another, the agent can call both simultaneously.

The entry points supported by agent framework are described in the following sections. With the exception of `monitor`, other entry points are optional. Each may be implemented in C++ or scripts.

Agent entry points described

This section describes each entry point in detail.

About the monitor entry point

The `monitor` entry point typically contains the logic to determine the status of a resource. For example, the `monitor` entry point of the IP agent checks whether or not an IP address is configured, and returns the state online, offline, or unknown.

Note: This entry point is mandatory.

The agent framework calls the `monitor` entry point after completing the `online` and `offline` entry points to determine if bringing the resource online or taking it offline was effective. The agent framework also calls this entry point periodically to detect if the resource was brought online or taken offline unexpectedly.

By default, the monitor entry point runs every sixty seconds (the default value of the `MonitorInterval` attribute) when a resource is online.

When a resource is expected to be offline, the monitor entry point does not run. The default value for the `OfflineMonitorInterval` attribute is 0.

The `monitor` entry point receives a resource name and `ArgList` attribute values as input (see “[ArgList reference attributes](#)” on page 165).

The entry point returns the resource status and the confidence level.

See “[Return values for entry points](#)” on page 36.

The entry point returns confidence level only when the resource status is online. The confidence level is informative only and is not used by the engine. It can be referenced by examining the value of `ConfidenceLevel` attribute.

A C++ entry point can return a confidence level of 0–100. A script entry point combines the status and the confidence level in a single number.

See “[Syntax for script entry points](#)” on page 97.

About the info entry point

The `info` entry point enables agents to obtain information about an online resource. For example, the Mount agent's `info` entry point could be used to report on space available in the file system. All information the `info` entry point collects is stored in the “temp” attribute `ResourceInfo`.

See “[About the ResourceInfo attribute](#)” on page 29.

The (script) entry point can optionally modify a resource's `ResourceInfo` attribute by adding or updating other name-value pairs using the following commands:

```
hares -modify res ResourceInfo -add key value
```

or

```
hares -modify res ResourceInfo -update key value
```

Refer to the `hares` manual page for more information on modifying values of string-association attributes.

See “[About the ResourceInfo attribute](#)” on page 29.

See “[Syntax for C++ entry points](#)” on page 49.

Return values for info entry point

- If the `info` entry point exits with 0 (success), the output captured on stdout for the script entry point, or the contents of the `info_output` argument for C++ entry point, is dumped to the `Msg` key of the `ResourceInfo` attribute. The `Msg` key is updated only when the `info` entry point is successful. The `State` key is set to the value: `Valid`.
- If the entry point exits with a non-zero value, `ResourceInfo` is updated to indicate the error; the script's stdout or the C++ entry point's `info_output` is ignored. The `State` key is set to the value: `Invalid`. The error message is written to the agent's log file.
- If the `info` entry point times out, output from the entry point is ignored. The `State` key is set to the value: `Invalid`. The error message is written to the agent's log file.
- If a user kills the `info` entry point (for example, `kill -15 pid`), the `State` key is set to the value: `Invalid`. The error message is written to the agent's log file.
See “[About logging agent messages](#)” on page 106.
- If the resource for which the entry point is invoked goes offline or faults, the `State` key is set to the value: `Stale`.
- If the `info` entry point is not implemented, the `State` key is set to the value: `Stale`. The error message is written to the agent's log file.

About the ResourceInfo attribute

The `ResourceInfo` attribute is a string association that stores name-value pairs. By default, there are three such name-value pairs:

- `State` indicates the status (valid, invalid, Stale) of the information contained in the `ResourceInfo` attribute.
- `Msg` indicates the output of the `info` entry point, if any.
- `TS` indicates the timestamp of when the `ResourceInfo` attribute was last updated.

These keys are updated only by the agent framework, not the entry point. The entry point can define and add other keys (name-value pairs) and update them.

The `ResourceInfo` (string-association) is a temporary attribute, the scope of which is set by the engine to be global for failover groups or local for parallel groups. Because `ResourceInfo` is a temporary attribute, its values are never dumped to the configuration file.

You can display the value of the `ResourceInfo` by using the `hares` command. The output of `hares -display` shows the first 20 characters of the current value; the output of `hares -value resourceResourceInfo` shows all name-value pairs in the keylist.

The resource for which the `info` entry point is invoked must be online.

When a resource goes offline or faults, the `State` key is marked “Stale” because the information is not current. If the `info` entry point exits abnormally, the `State` key is marked “Invalid” because not all of the information is known to be valid. Other key data, including `Msg` and `TS` keys, are not touched. You can manually clear values of the `ResourceInfo` attribute by using the `hares -flushinfo` command. This command deletes any optional keys for the `ResourceInfo` attribute and sets the three mandatory keys to their default values.

Invoking the info entry point

You can invoke the `info` entry point from the command line for a given online resource using the `hares -refreshinfo` command.

By setting the `InfoInterval` attribute to some value other than 0, you can configure the agent to invoke the `info` entry point periodically for an online resource.

See “[InfoInterval](#)” on page 169.

About the online entry point

The `online` entry point typically contains the code to bring a resource online. For example, the `online` entry point for an IP agent configures an IP address. When the online procedure completes, the `monitor` entry point is automatically called by the framework to verify that the resource is online.

The `online` entry point receives a resource name and `ArgList` attribute values as input. It returns an integer indicating the number of seconds to wait for the online to take effect. The typical return value is 0. If the return value is not zero, the agent framework waits the number of seconds indicated by the return value before calling the `monitor` entry point for the resource.

About the offline entry point

The `offline` entry point takes a resource offline. For example, the `offline` entry point for an IP agent removes an IP address from the system. When the offline procedure completes, the `monitor` entry point is automatically called by the framework to verify that the resource is offline.

The `offline` entry point receives a resource name and `ArgList` attribute values as input. It returns an integer indicating the number of seconds to wait for the offline to take effect. The typical return value is 0. If the return value is not zero, the agent framework waits the number of seconds indicated by the return value to call the `monitor` entry point for the resource.

About the clean entry point

The `clean` entry point is called by the agent framework when all ongoing tasks associated with a resource must be terminated and the resource must be taken offline, perhaps forcibly. The entry point receives as input the resource name, an encoded reason describing why the entry point is being called, and the `ArgList` attribute values. It must return 0 if the operation is successful and 1 if unsuccessful.

The reason for calling the entry point is encoded according to the following enum type:

```
enum VCSAgWhyClean {  
    VCSAgCleanOfflineHung,  
    VCSAgCleanOfflineIneffective,  
    VCSAgCleanOnlineHung,  
    VCSAgCleanOnlineIneffective,  
    VCSAgCleanUnexpectedOffline,  
    VCSAgCleanMonitorHung  
};
```

For script-based Clean entry points, the Clean reason is passed as an integer:

```
0 => offline hung  
1 => offline ineffective  
2 => online hung  
3 => online ineffective  
4 => unexpected offline  
5 => monitor hung
```

The above is an enum type, so same integer value is passed irrespective of whether the entry point is written in C++ or is script-based.

Note: On Windows, `VCSAgCleanOfflineHung`, `VCSAgCleanOnlineHung`, and `VCSAgCleanMonitorHung` do not apply. Entry point timeouts are not configurable parameters on the Windows platform.

- **VCSAgCleanOfflineHung**
The `offline` entry point did not complete within the expected time.
This parameter is not used on Windows.
See “[OfflineTimeout](#)” on page 173.
- **VCSAgCleanOfflineIneffective**
The `offline` entry point was ineffective. The `monitor` entry point returned a status other than `OFFLINE` after the scheduled invocation of the `offline` entry point for the resource.
- **VCSAgCleanOnlineHung**
The `online` entry point did not complete within the expected time.
This parameter is not used on Windows.

(See “[OnlineTimeout](#)” on page 174.)

- **VCSAgCleanOnlineIneffective**
The `online` entry point was ineffective. The `monitor` entry point scheduled for the resource after the `online` entry point invocation returned a status other than `ONLINE`.
- **VCSAgCleanUnexpectedOffline**
The online resource faulted because it was taken offline unexpectedly.
- **VCSAgCleanMonitorHung**
The online resource faulted because the `monitor` entry point consistently failed to complete within the expected time.
This parameter is not used on Windows.
(See “[FaultOnMonitorTimeouts](#)” on page 169.)

The agent supports the following tasks when the `clean` entry point is implemented:

- Automatically restarts a resource on the local system when the resource faults.
See “[RestartLimit](#)” on page 177.
- Automatically retries the `online` entry point when the attempt to bring a resource online fails.
See “[OnlineRetryLimit](#)” on page 174.
- Enables the engine to bring a resource online on another system when the `online` entry point for the resource fails on the local system.

For the above actions to occur, the `clean` entry point must run successfully, that is, return an exit code of 0.

About the action entry point

Runs a prespecified action on a resource. Use the entry point to run non-periodic actions like suspending a database or resuming the suspended database.

The `SupportedActions` attribute is a keylist attribute that lists all the actions that are intended on being supported. Each action is identified by name, or `action_token`.

See “[SupportedActions](#)” on page 179.

For an agent, all action entry points must be either C++ or script-based; you cannot use both C++ and scripts.

Make sure the action scripts reside within an `actions` directory under the agent directory. Create a script for each action. Use the correct `action_token` as the script name.

For example, a script called `suspend` defines the actions to be performed when the `action_token` “suspend” is invoked via the `hares -action` command.

For C++ entry points, actions are implemented via a switch statement that defines a case for each possible `action_token`.

See “[Syntax for C++ action](#)” on page 58.

The following shows the syntax for the `-action` option used with the `hares` command:

```
hares -action res_name action_token [-actionargs arg1 arg2 ... ]
      [-sys sys_name] [-user user@domain] [-domain domaintype]
```

The following example commands show the invocation of the action entry point using the example action tokens, `DBSuspend` and `DBResume`:

```
hares -action DBResource DBSuspend -actionargs dbsuspend -sys
Sys1
```

Also,

```
hares -action DBResource DBResume -actionargs dbstart -sys Sys1
```

See also “[RegList](#)” on page 175.

Return values for action entry point

The action entry point exits with a 0 if it is successful, or 1 if not successful.

The command `hares -action` exits with 0 if the action entry point exits with a 0 and 1 if the action entry point is not successful.

The agent framework limits the output of the script-based action entry point to 2048 bytes.

Output refers to information that the script prints to `stdout` or `stderr`. When users run the `hares -action` command, the command prints this output. The output is also logged to the `vcstoneclientd` log file.

About the attr_changed entry point

This entry point provides a way to respond to resource attribute value changes. The `attr_changed` entry point is called when a resource attribute is modified, and only if that resource attribute is registered with the agent framework for notification.

See “[VCSAgRegister](#)” on page 67.

See “[VCSAgUnregister](#)” on page 68.

Registering can be accomplished either through the `VCSAgRegister` method for C++ based agents or by definition in the `RegList` for script based agents.

See “[RegList](#)” on page 175.

The `attr_changed` entry point receives as input the resource name registered with the agent framework for notification, the name of the changed resource, the name of the changed attribute, and the new attribute value. It does not return a value.

About the open entry point

When an agent starts, the `open` entry point of each configured resource is called before its `online`, `offline`, or `monitor` entry points are called. This allows you to include initialization for specific resources. Most agents do not require this functionality and will not implement this entry point.

The `open` entry point is also called whenever the `Enabled` attribute for the resource changes from 0 to 1. The entry point receives the resource name and `ArgList` attribute values as input and returns no value.

A resource can be brought online, taken offline, and monitored only if it is managed by an agent. For an agent to manage a resource, the value of the resource’s `Enabled` attribute must be set to 1. The entry point creates the environment needed for other entry points to function. For example, the entry point could create files required by other entry points for the resource, or perform some resource-specific setup.
points for the resource, or perform some resource-specific setup.

About the close entry point

The `close` entry point is called whenever the `Enabled` attribute for a resource changes from 1 to 0, or when a resource is deleted from the configuration on a running cluster and the state of the resource permits running the close entry point.

Note that a resource is monitored only if it is managed by an agent. For an agent to manage a resource, the resource's `Enabled` attribute value must be set to 1.

See the table below to find out which states of the resource allow running of the close entry point when the resource is deleted on a running cluster. It receives a resource name and `ArgList` attribute values as input and returns no value. This entry point typically deinitializes the resource if implemented. Most agents do not require this functionality and will not implement this entry point.

Table 2-1 States in which CLOSE entry point runs - based on operations type of resource

Resource Type	Online State	Offline State	Probing	Going Offline Waiting	Going Online Waiting
None (persistent)	Yes	N/A	Yes	Yes	N/A
OnOnly	Yes	Yes	Yes	Yes	Yes
OnOff	No	Yes	Yes	No	No

The open and close entry points are related in the sense that the open entry point creates the environment needed for other entry points, while the close entry points cleans the setup created by open entry point.

About the shutdown entry point

The `shutdown` entry point is called before the agent shuts down. It performs any agent cleanup required before the agent exits. It receives no input and returns no value. Most agents do not require this functionality and do not implement this entry point.

Return values for entry points

The following table summarizes the return values for each entry point.

Table 2-2 Return values for entry points

Entry Point	Return Values
Monitor	C++ Based Returns ResStateValues: ■ VCSAgResOnline ■ VCSAgResOffline ■ VCSAgResUnknown ■ VCSAgResIntentionalOffline Script-Based Exit values: ■ 99 - Unknown ■ 100 - Offline ■ 101-110 - Online ■ 200 - Intentional Offline ■ Other values - Unknown.
Info	0 if successful; non-zero value if not successful
Online	Integer specifying number of seconds to wait before monitor can check the state of the resource; typically 0, that is, check resource state immediately.
Offline	Integer specifying number of seconds to wait before monitor can check the state of the resource; typically 0, that is, check resource state immediately.
Clean	0 if successful; non-zero value if not successful If clean fails, the resource remains in a transition state awaiting the next periodic monitor. After the periodic monitor, clean is attempted again. The sequence of clean attempt followed by monitoring continues until clean succeeds. See “CleanRetryLimit” on page 166. See Chapter 9, “State transition diagrams” on page 181 for descriptions of internal transition states.
Action	0 if successful; non-zero value if not successful
Attr_changed	None
Open	None
Close	None
Shutdown	None

Considerations for using C++ or script entry points

You may implement an entry point as a C++ function or a script.

- The advantage to using C++ is that entry points are compiled and linked with the agent framework library. They run as part of the agent process, so no system overhead for creating a new process is required when they are called. Also, since the entry point invocation is just a function call, the execution of the entry point is relatively faster. However, if the functionality of an entry point needs to be changed, the agent would need to be recompiled to make the changes take effect.
- The advantage to using scripts is that you can modify the entry points dynamically. However, to run the script, a new process is created for each entry point invocation, so the execution of an entry point is relatively slower and uses more system resource compared to the C++ implementation.

Note that you may use C++ or scripts in any combination to implement multiple entry points for a single agent. This allows you to implement each entry point in the most advantageous manner. For example, you may use scripts to implement most entry points while using C++ to implement the `monitor` entry point, which is called often. If the `monitor` entry point were written in script, the agent must create a new process to run the monitor entry point each time it is called.

See [Chapter 3, “Creating entry points in C++”](#) on page 45.

See [Chapter 4, “Creating entry points in scripts”](#) on page 93.

About the VCSAgStartup routine

When an agent starts, it uses the routine named `VCSAgStartup` to initialize the agent's data structures.

If you implement entry points using scripts

If you implement all of the agent's entry points as scripts:

On UNIX, use the `Script51Agent` binary.

On Windows, use the `DefaultAgent.dll` or the `DefaultAgent50Agent.dll`.

The built-in implementation of `VCSAgStartup()` in these binaries initializes the agent's data structures such that it causes the agent to look for and execute the scripts for the entry points.

See [Chapter 4, “Creating entry points in scripts”](#) on page 93.

If you implement all or some of the entry points in C++

If you develop an agent with at least one entry point implemented in C++, you must implement the function `VCSAgStartup()` and use the required C++ primitives to register the C++ entry point with the agent framework.

Example: `VCSAgStartup` with C++ and script entry points

When using C++ to implement an entry point, use the `VCSAgInit` API and specify the entry point and the function name. In the following example, the function `my_shutdown` is defined as the Shutdown entry point.

```
#include "VCSAgApi.h"
void my_shutdown() {
    ...
}

void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(v51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPShutdown, my_shutdown);
}
```

Note that the `monitor` entry point, which is mandatory, is assigned a `NULL` value, indicating it is implemented using scripts. If you are using a script entry point, or if you are not implementing an optional entry point, set the corresponding field to `NULL`.

For an entry point whose field is set to `NULL`, the agent automatically looks for the correct script to execute:

UNIX: `$VCSONE_HOME/bin/resource_type/<entry_point>`

Windows:

`%VCSONE_HOME%\bin\resource_type\<entry_point>.extn`

Note: For Windows, the extension (.extn) assumes the value “exe” for executable programs, “sh” for shell scripts compatible with MKS Toolkit, “bat” or “cmd” for batch files, “pl” for Perl scripts, “js” for Javascript, or “vbs” for Visual Basic scripts.

About the ArgList and ArgListValues attributes

The ArgList attribute specifies which attributes need to be passed to agent entry points. The agent framework populates the ArgListValues attribute with the list of attributes and their associated values.

In C++ agents, the value of the ArgListValues attribute is passed through a parameter of type `void **`. For example, the signature of the online entry point is:

```
unsigned int  
res_online(const char *res_name, void **attr_val);
```

In script agents, the value of the ArgListValues attribute is passed as command-line arguments to the entry point script.

ArgListValues attribute for agents registered as V50 and later

For agents registered as V50 or later, the ArgListValues attribute specifies the attributes and their values in tuple format.

- For scalar attributes, there are three components that define the ArgListValues attribute.
 - The name of the attribute
 - The number of elements in the value, which for scalar attributes is always 1
 - The value itself
- For non-scalar attributes (vector, keylist, and association), for each attribute there are N+2 components in the ArgListValues attribute, where N equals the number of elements in the attribute's value.
 - The name of the attribute
 - The number of elements in the attribute's value
 - The remaining N elements correspond to the attribute's value. Note that N could be zero.

Overview of the name-value tuple format

Pre-5.0 agents required that the arguments passed to the entry point to be in the order indicated by the ArgList attribute as it was defined in the resource type. The order of parsing the arguments was determined by their position the resource type definition.

With the agent framework for V50 and later, agents can use entry points that can be passed attributes and their values in a format of name-value tuples. Such a format means that attributes and their values are parsed by the name of the attribute and not by their position in the ArgList Attribute.

The general tuple format for attributes in the ArgList is:

<name> <number_of_elements_in_value> <value>

Scalar attribute format

For *scalar* attributes, whether string, integer, or boolean, the formatting is:

<attribute_name> 1 <value>

Example is:

DiskGroupName 1 mydg

Vector attribute format

For *vector* attributes, whether string or integer, the formatting is:

<attribute_name> <number_of_values_in_vector> <values_in_vector>

Examples are:

MyVector 3 aa cc dd

MyEmptyVector 0

Keylist attribute format

For string *keylist* attributes, the formatting is:

<attribute_name> <number_of_keys_in_keylist> <keys>

Examples are:

DiskAttr 4 hdisk3 hdisk4 hdisk5 hdisk6

DiskAttr 0

Association attribute format

For association attributes, whether string or integer, the formatting is:

<attribute_name> <number_of_keys_and_values> <values_of_keylist>

Examples are:

MyAssoc 4 key1 val1 key2 val2

MyAssoc 0

About the entry point timeouts

Use the AEPTIMEOUT attribute to append the timeout value for a particular entry the list of arguments passed to the entry point.

This feature does not apply to pre-V50 agents.

If you set AEPTIMEOUT to 1, the agent framework passes the timeout value for an entry point as an argument for the entry point in the name-value tuple format.

The name of the attribute that gets passed is called AEPTIMEOUT.

This makes the task of retrieving information about entry point timeout values easy for agent developers. Instead of looking for different strings like `MonitorTimeout` and `CleanTimeout`, agent developers just need to look for the string `AEPTIMEOUT`.

For example, if an agent uses an attribute called `PathName` set to `/tmp/foo`, the parameters passed to the monitor entry point are:

If AEPTimeout is set to 0 <resource-name> PathName 1 /tmp/foo.

If AEPTIMEOUT set to 1 <resource-name> PathName 1 /tmp/foo
AEPTIMEOUT 1 <value of MonitorTimeout attribute>

Applying the same example for the clean entry point, the parameters are:

If AEPTimeout is set to 0 <resource-name> <clean reason> PathName 1 /tmp/
foo

If AEPTIMEOUT is set to 1 <resource-name> <clean reason> PathName 1 /tmp/
foo AEPTIMEOUT 1 <value of CleanTimeout
attribute>

If the timeout attribute is overridden at the resource level, this mechanism takes care of passing the overridden value to the entry points for that resource.

See “[AEPTimeout](#)” on page 164.

ArgListValues attribute for agents registered as V40 and earlier

For agents registered as V40 and earlier, the ArgListValues attribute is an ordered list of attribute values. The attribute values are listed in the same order as in the ArgList attribute.

For example, if Type “Foo” is defined in the file types.xml as:

```
type Foo (
  <attributes class="static">
    <attribute name="ArgList" type="str" dimension="vector">
      <default>
        <val>IntAttr</val>
        <val>StringAttr</val>
        <val>VectorAttr</val>
        <val>AssocAttr</val>
      </default>
    </attribute>
  </attributes>
  <attributes>
    <attribute name="IntAttr" type="int" dimension="scalar">
      <display_name>"IntAttr"</display_name>
      <must_configure>1</must_configure>
    </attribute>
  </attributes>
)
```

```
        <unique>1</unique>
        <important>1</important>
    </attribute>
    <attribute name="StringAttr" type="str"
        dimension="scalar">
        <display_name>"StringAttr"</display_name>
        <must_configure>1</must_configure>
        <unique>1</unique>
        <important>1</important>
    </attribute>
    <attribute name="VectorAttr" type="str"
        dimension="vector">
        <display_name>"VectorAttr"</display_name>
        <must_configure>1</must_configure>
        <unique>1</unique>
        <important>1</important>
    </attribute>
    <attribute name="AssocAttr" type="int"
        dimension="assoc">
        <display_name>"AssocAttr"</display_name>
        <must_configure>1</must_configure>
        <unique>1</unique>
        <important>1</important>
    </attribute>
</attributes>
)
```

And the resource “Bar” is defined in the main.xml file as:

```
<resources>
  <resource name="Bar" type="Foo">
    <attributes>
      <attribute name="IntAttr"><scalar>100</scalar>
      </attribute>
      <attribute name="StringAttr">
        <scalar>"Oracle"</scalar>
      </attribute>
      <attribute name="VectorAttr"><vector>
        <val>"vol1"</val>
        <val>"vol2"</val>
        <val>"vol3"</val></vector>
      </attribute>
      <attribute name="AssocAttr"><assoc>
        <val key="disk1">1024</val>
        <val key="disk2">512</val></assoc>
      </attribute>
    </attributes>
  </resource>
</resources>
```

Then, for V50 and later, the parameter attr_val is:

```
attr_val[0] = "IntAttr"
attr_val[1] = "1"      // Number of components in
                      // IntAttr attr value
attr_val[2] = "100"   // Value of IntAttr
attr_val[3] = "StringAttr"
attr_val[4] = "1"      // Number of components in
                      // StringAttr attr value
attr_val[5] = "Oracle" // Value of StringAttr
attr_val[6] = "VectorAttr"
attr_val[7] = "3"      // Number of components in
                      // VectorAttr attr value

attr_val[8] = "vol1"
attr_val[9] = "vol2"
attr_val[10] = "vol3"
attr_val[11] = "AssocAttr"
attr_val[12] = "4"     // Number of components in
                      // AssocAttr attr value

attr_val[13] = "disk1"
attr_val[14] = "1024"
attr_val[15] = "disk2"
attr_val[16] = "512"
attr_val[17] = NULL    // Last element
```

Or, for V40 and earlier, the parameter attr_val is:

```
attr_val[0] ==>  "100"      // Value of IntAttr, the first
                        // ArgList attribute.
attr_val[1] ==>  "Oracle"   // Value of StringAttr.
attr_val[2] ==>  "3"        // Number of components in
                        // VectorAttr.

attr_val[3] ==>  "vol1"
attr_val[4] ==>  "vol2"
attr_val[5] ==>  "vol3"
attr_val[6] ==>  "4"        // Number of components in
                        // AssocAttr.

attr_val[7] ==>  "disk1"
attr_val[8] ==>  "1024"
attr_val[9] ==>  "disk2"
attr_val[10] ==> "512"
attr_val[11] ==> NULL      // Last element.
```

Creating entry points in C++

- [About creating entry points in C++](#)
- [Data Structures](#)
- [Syntax for C++ entry points](#)
- [Agent framework primitives](#)
- [Agent Framework primitives for container support](#)
- [Agent framework primitives for Windows](#)

About creating entry points in C++

Because the agent framework is multithreaded, all C++ code written by the agent developer must be MT-safe. For best results, avoid using global variables. If you do use them, access must be serialized (for example, by using mutex locks).

The following guidelines also apply:

- Do not use C library functions that are unsafe in multithreaded applications. Instead, use the equivalent reentrant versions, such as `readdir_r()` instead of `readdir()`. Access manual pages for either of these commands by entering: `man command`.
- When acquiring resources (dynamically allocating memory or opening a file, for example), use thread-cancellation handlers to ensure that resources are freed properly. See the manual pages for `pthread_cleanup_push` and `pthread_cleanup_pop` for details. Access manual pages for either of these commands by entering: `man command`.

If you develop an agent with at least one entry point implemented in C++, you must implement the function `VCSAgStartup()` and use the required C++ primitives to register the C++ entry point with the agent framework.

A sample file containing templates for creating an agent using C++ entry points is located in:

UNIX: `$VCSONE_HOME/src/agent/Sample`

Windows: `%VCSONE_HOME\src\agent\Sample`

You can use C++ to develop agents for monitoring applications that run in containers, including non-global zones. VCS One provides APIs for container support.

See “[Agent Framework primitives for container support](#)” on page 82.

Entry point examples in this chapter

In this chapter, the example entry points are shown for an agent named Foo. The example agent has the following resource type definition:

Entry Point	What it does in this agent
online	Creates a file as specified by the Pathname attribute
monitor	Checks for the existence of a file specified by the PathName attribute
offline	Deletes the file specified by the PathName attribute
clean	Forcibly deletes the file specified by the PathName attribute
action	Runs a prespecified action
info	Populates the ResourceInfo attribute with the values of the attributes specified by the PathName attribute

Data Structures

The VCSAgResState enumeration supports the following return codes from the monitor entry point.

- VCSAgResOffline
- VCSAgResOnline
- VCSAgResUnknown
- VCSAgResIntentionalOffline (only in V51 and later agents)

```
// Values for the reason why the clean entry point
// is called.

enum VCSAgWhyClean {
    VCSAgCleanOfflineHung,    // offline entry point did
                              // not complete within the
                              // expected time.
    VCSAgCleanOfflineIneffective, // offline entry point
                              // was ineffective.
    VCSAgCleanOnlineHung,     // online entry point did
                              // not complete within the
                              // expected time.
    VCSAgCleanOnlineIneffective, // online entry point
                              // was ineffective.
    VCSAgCleanUnexpectedOffline, // the resource became
                              // offline unexpectedly.
    VCSAgCleanMonitorHung     // monitor entry point did
                              // not complete within the
                              // expected time.
};
```

Note: On Windows, VCSAgCleanOfflineHung, VCSAgCleanOnlineHung, and VCSAgCleanMonitorHung do not apply. Entry point timeouts are not configurable parameters on Windows.

Syntax for C++ entry points

This section describes the syntax for C++ entry points.

Syntax for C++ VCSAgStartup

```
void VCSAgStartup();
```

Note that the name of the C++ function must be `VCSAgStartup()`.

For example:

```
// This example shows the VCSAgStartup() entry point
// implementation, assuming that the monitor, online, offline
// and clean entry points are implemented in C++ and the
// respective function names are res_monitor, res_online,
// res_offline, and res_clean.

#include "VCSAgApi.h"
void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}

VCSAgResState res_monitor(const char *res_name, void
                           **attr_val, int
                           *conf_level) {
    ...
}

unsigned int res_online(const char *res_name,
                       void **attr_val) {
    ...
}

unsigned int res_offline(const char *res_name,
                        void **attr_val) {
    ...
}
```

Syntax for C++ monitor

```
VCSAgResState  
res_monitor(const char *res_name, void **attr_val, int  
*conf_level);
```

You may select any name for the function.

The parameter `conf_level` is an output parameter. The return value, which indicates the resource status, must be a defined `VCSAgResState` value.

See [“Return values for entry points”](#) on page 36.

For example:

```
#include "VCSAgApi.h"  
  
VCSAgResState  
res_monitor(const char *res_name, void **attr_val, int  
*conf_level)  
{  
  
    // Code to determine the state of a resource.  
    VCSAgResState res_state = ...  
    if (res_state == VCSAgResOnline) {  
        // Determine the confidence level (0 to 100).  
        *conf_level = ...  
    }  
    else {  
        *conf_level = 0;  
    }  
    return res_state;  
}  
  
void VCSAgStartup()  
{  
    VCSAG_LOG_INIT("VCSAgStartup");  
  
    VCSAgSetLogCategory(10051);  
    VCSAgInitEntryPointStruct(V51);  
  
    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);  
}
```

Syntax for C++ info

```
unsigned int (*info) (const char *res_name,
                      VCSAgResInfoOp resinfo_op, void **attr_val, char
                      **info_output, char ***opt_update_args, char
                      ***opt_add_args);
```

You may select any name for the function.

resinfo_op

The `resinfo_op` parameter indicates whether to initialize or update the data in the `ResourceInfo` attribute. The values of this field and their significance are described in the following table:

Value of resinfo_op	Significance
1	Add non-default keys to the three default keys State, Msg, and TS and initialize the name-value data pairs in the <code>ResourceInfo</code> attribute. This invocation indicates to the entry point that the current value of the <code>ResourceInfo</code> attribute contains only the basic three keys State, Msg, and TS.
2	Update only the non-default key-value data pairs in the <code>ResourceInfo</code> attribute, not the default keys State, Msg, and TS. This invocation indicates that <code>ResourceInfo</code> attribute contains non-default keys in addition to the default keys and only the non-default keys are to be updated. Attempt to add keys with this invocation will result in errors.

info_output

The parameter `info_output` is a character string that stores the output of the `info` entry point. The output value could be any summarized data for the resource. The `Msg` key in the `ResourceInfo` attribute is updated with `info_output`. If the `info` entry point exits with success (0), the output stored in `info_output` is dumped into the `Msg` key of the `ResourceInfo` attribute.

The `info` entry point is responsible for allocating memory for `info_output`. The agent framework handles the deletion of any memory allocated to this argument. Since memory is allocated in the entry point and deleted in the agent framework, the entry point needs to pass the address of the allocated memory to the agent framework.

opt_update_args

The `opt_update_args` parameter is an array of character strings that represents the various name-value pairs in the `ResourceInfo` attribute. This argument is allocated memory in the `info` entry point, but the memory allocated for it will be freed in the agent framework. The `ResourceInfo` attribute is updated with these name-value pairs. The names in this array must already be present in the `ResourceInfo` attribute.

For example:

```
ResourceInfo = { State = Valid, Msg = "Info entry point output",  
                TS = "Wed May 28 10:34:11 2003", FileOwner = root,  
                FileGroup = root, FileSize = 100 }
```

A valid `opt_update_args` array for this `ResourceInfo` attribute would be:

```
opt_update_args = { "FileSize", "102" }
```

This array of name-value pairs updates the dynamic data stored in the `ResourceInfo` attribute.

An invalid `opt_update_args` array would be one that specifies a key not already present in the `ResourceInfo` attribute or one that specifies any of the keys: `State`, `Msg`, or `TS`. These three keys can only be updated by the agent framework and not by the entry point.

opt_add_args

`opt_add_args` is an array of character strings that represent the various name-value pairs to be added to the `ResourceInfo` attribute. The names in this array represent keys that are *not* already present in the `ResourceInfo` association list and have to be added to the attribute. This argument is allocated memory in the `info` entry point, but this memory is freed in the agent framework. The `ResourceInfo` attribute is populated with these name-value pairs.

For example:

```
ResourceInfo = { State = Valid, Msg = "Info entry point output",  
                TS = "Wed May 28 10:34:11 2003" }
```

A valid `opt_add_args` array for this would be:

```
opt_add_args = { "FileOwner", "root", "FileGroup",  
                "root",  
                "FileSize", "100" }
```

This array of name-value pairs adds to and initializes the static and dynamic data stored in the `ResourceInfo` attribute.

An invalid `opt_add_args` array would be one that specifies a key that is already present in the `ResourceInfo` attribute, or one that specifies any of the keys `State`, `Msg`, or `TS`; these are keys that can be updated only by the agent framework, not by the entry point.

Example: info entry point implementation in C++

Set the `VCSAgValidateAndSetEntryPoint()` parameter to the name of the entry point's function (`res_info`).

Allocate the `info` output buffer in the entry point as shown in the example below. The buffer can be any size (the example uses 80), but the agent framework truncates it to 2048 bytes. For the optional name-value pairs, name and value each have a limit of 4096 bytes (the example uses 15).

Example V51 entry point:

```

extern "C" unsigned int res_info(const char *res_name,
VCSAgResInfoOp resinfo_op, void **attr_val, char **info_output,
char ***opt_update_args, char ***opt_add_args)
{
    struct stat stat_buf;
    int i;
    char **args = NULL;
    char *out = new char [80];

    *info_output = out;

    VCSAgSnprintf(out, 80, "Output of info entry point - updates
        the \"Msg\" key in ResourceInfo attribute");

    // Use the stat system call on the file to get its
    // information The attr_val array will look like "PathName"
    // "1" "<pathname value>" ... Assuming that PathName is the
    // first attribute in the attr_val array, the value
    // of this attribute will be in index 2 of this attr_val
    // array

    if (attr_val[2]) {

        if ((strlen((CHAR *) (attr_val[2])) != 0) &&
            (stat((CHAR *) (attr_val[2]), &stat_buf) == 0)) {

            if (resinfo_op == VCSAgResInfoAdd) {
                // Add and initialize all the static and
                // dynamic keys in the ResourceInfo attribute
                args = new char * [7];
                for (i = 0; i < 6; i++) {
                    args[i] = new char [15];
                }

                // All the static information - file owner
                // and group
                VCSAgSnprintf(args[0], 15, "%s", "Owner");
                VCSAgSnprintf(args[1], 15, "%d",
                    stat_buf.st_uid);
                VCSAgSnprintf(args[2], 15, "%s", "Group");
            }
        }
    }
}

```

```
        VCSAgSnprintf(args[3], 15, "%d",
stat_buf.st_gid);

        // Initialize the dynamic information for the file
        VCSAgSnprintf(args[4], 15, "%s", "FileSize");
        VCSAgSnprintf(args[5], 15, "%d",
stat_buf.st_size);
        args[6] = NULL;
        *opt_add_args = args;
    }
    else {

        // Simply update the dynamic keys in the
        // ResourceInfo attribute. In this case, the
        // dynamic info on the file
        args = new char * [3];
        for (i = 0; i < 2; i++) {
            args[i] = new char [15];
        }
        VCSAgSnprintf(args[0], 15, "%s", "FileSize");
        VCSAgSnprintf(args[1], 15, "%d",
stat_buf.st_size);
        args[2] = NULL;
        *opt_update_args = args;
    }
}
else {
    // Set the output to indicate the error
    VCSAgSnprintf(out, 80, "Stat on the file %s failed",
attr_val[2]);
    return 1;
}
}
else {
    // Set the output to indicate the error
    VCSAgSnprintf(out, 80, "Error in arglist values passed to
the info entry point");
    return 1;
}

// Successful completion of the info entry point
return 0;

} // End of entry point definition
```

Syntax for C++ online

```
unsigned int  
res_online(const char *res_name, void **attr_val);
```

You may select any name for the function.

Set the `VCSAgValidateAndSetEntryPoint()` parameter to the name of the entry point's function.

In the following example, the function `res_online` is defined as the Online entry point.

For example:

```
#include "VCSAgApi.h"  
  
unsigned int  
res_online(const char *res_name, void **attr_val) {  
    // Implement the code to online a resource here.  
    ...  
    // If monitor can check the state of the resource  
    // immediately, return 0. Otherwise, return the  
    // appropriate number of seconds to wait before  
    // calling monitor.  
    return 0;  
}  
  
void VCSAgStartup()  
{  
    VCSAG_LOG_INIT("VCSAgStartup");  
  
    VCSAgSetLogCategory(10051);  
    VCSAgInitEntryPointStruct(V51);  
  
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);  
}
```

Syntax for C++ offline

```
unsigned int  
res_offline(const char *res_name, void **attr_val);
```

You may select any name for the function.

Set the `VCSAgValidateAndSetEntryPoint()` parameter to the name of the entry point's function.

In the following example, the function `res_offline` is defined as the Offline entry point.

For example:

```
#include "VCSAgApi.h"  
  
unsigned int  
res_offline(const char *res_name, void **attr_val) {  
    // Implement the code to offline a resource here.  
    ...  
    // If monitor can check the state of the resource  
    // immediately, return 0. Otherwise, return the  
    // appropriate number of seconds to wait before  
    // calling monitor.  
    return 0;  
}  
  
void VCSAgStartup()  
{  
    VCSAG_LOG_INIT("VCSAgStartup");  
  
    VCSAgSetLogCategory(10051);  
    VCSAgInitEntryPointStruct(V51);  
  
    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);  
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);  
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);  
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);  
}
```

Syntax for C++ clean

```
unsigned int  
res_clean(const char *res_name, VCSAgWhyClean reason, void  
**attr_val);
```

You may select any name for the function.

Set the `VCSAgValidateAndSetEntryPoint()` parameter to the name of the entry point's function.

In the following example, the function `res_clean` is defined as the Clean entry point.

For example:

```
#include "VCSAgApi.h"  
  
unsigned int  
res_clean(const char *res_name, VCSAgWhyClean reason,  
          void **attr_val) {  
    // Code to forcibly offline a resource.  
    ...  
    // If the procedure is successful, return 0; else  
    // return 1.  
    return 0;  
  
void VCSAgStartup()  
{  
    VCSAG_LOG_INIT("VCSAgStartup");  
  
    VCSAgSetLogCategory(10051);  
    VCSAgInitEntryPointStruct(V51);  
  
    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);  
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);  
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);  
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);  
}
```

Syntax for C++ action

```
unsigned int
action(const char *res_name, const char *action_token,
       void **attr_val, char **args, char *action_output);
```

The parameters passed to the C++ action entry point are described as follows using the example that the user fires

\$> hares -action res1 myaction...

from the command-line or the equivalent from the GUI.

- **res_name:** This is an input parameter. The name of the resource in whose context the action entry point is being invoked. In the above example, res_name would be set to "res1".
- **action_token:** This is an input parameter. This gives the name of the action that the user wants to run. In the above example, action_token would be set to "myaction".

If the user ran

\$> hares -action res1 youraction ...

then the same function above will get invoked but action_token will be set to "youraction". This parameter enables different actions to be implemented for the same agent which will all get handled in the same function above.

- **attr_val:** This is an input parameter. This contains the ArgListValues of the resource for which the action is invoked.
- **args:** This is an input parameter. This contains the list of strings that are passed to the "-actionargs" switch when invoking the "hares -action" command.

\$> hares -action res1 myaction -actionargs foo bar fubar -sys ...

would give "foo", "bar" and "fubar" in the args parameter.

- **action_output:** This is an output parameter. Any output that the agent developer wants the user to see as a result of invoking the "hares -action" command needs to be filled into the buffer whose pointer is given by this parameter. The maximum number of characters that will be displayed to the user is 2048 (2K).

Use the VCSAgValidateAndSetEntryPoint() API to register the name of the function that implements the action entry-point for the agent.

For example:

```
extern "C"
unsigned int res_action (const char *res_name, const char
                        *token,void **attr_val, char **args, char
                        *action_output)
```

```

    {
const int output_buffer_size = 2048;
    //
    // checks on the attr_val entry point arg list
    // perform an action based on the action token passed in

if (!strcmp(token, "token1")) {
    //
    // Perform action corresponding to token1
    //
} else if (!strcmp(token, "token2")) {
    //
    // Perform action corresponding to token2
    //
}
:
:
:
} else {
    //
    // a token for which no action is implemented yet
    //
VCSAgSprintf(action_output, output_buffer_size, "No implementation
provided for token(%s)", token);
}

    //
    // Any other checks to be done
    //
    //
    // return value should indicate whether the ep succeeded or
    // not:
    // return 0 on success
    // any other value on failure
    //
if (success) {
return 0;
}
else {
return 1;
}
}
    
```

Syntax for C++ attr_changed

```
void  
res_attr_changed(const char *res_name, const char  
                *changed_res_name,  
                const char *changed_attr_name,  
                void **new_val);
```

The parameter `new_val` contains the attribute's new value. The encoding of `new_val` is similar to the encoding of the “[About the ArgList and ArgListValues attributes](#)” on page 39.

You may select any name for the function.

Set the `VCSAgValidateAndSetEntryPoint()` parameter to the name of the entry point's function.

In the following example, the function `res_attr_changed` is defined as the `attr_changed` entry point.

Note: This entry point is called only if you register for change notification using the primitive “[VCSAgRegister](#)” on page 67, or the agent parameter `RegList` (see “[RegList](#)” on page 175).

For example:

```
#include "VCSAgApi.h"  
  
void  
res_attr_changed(const char *res_name,  
                const char *changed_res_name,  
                const char *changed_attr_name,  
                void **new_val) {  
    // When the value of attribute Foo changes, take some action.  
    if ((strcmp(res_name, changed_res_name) == 0) &&  
        (strcmp(changed_attr_name, "Foo") == 0)) {  
        // Extract the new value of Foo. Here, it is assumed  
        // to be a string.  
        const char *foo_val = (char *)new_val[0];  
        // Implement the action.  
        ...  
    }  
}
```

```

// Resource Oral managed by this agent needs to
// take some action when the Size attribute of
// the resource Disk1 is changed.
if ((strcmp(res_name, "Oral") == 0) &&
    (strcmp(changed_attr_name, "Size") == 0) &&
    (strcmp(changed_res_name, "Disk1") == 0)) {

    // Extract the new value of Size. Here, it is
    // assumed to be an integer.
    int sizeval = atoi((char *)new_val[0]);
    // Implement the action.
    ...
}
}

void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}

```

Syntax for C++ open

```
void res_open(const char *res_name, void **attr_val);
```

You may select any name for the function.

Set the VCSAgValidateAndSetEntryPoint() parameter to the name of the entry point's function.

In the following example, the function res_open is defined as the Open entry point.

For example:

```
#include "VCSAgApi.h"

void res_open(const char *res_name, void **attr_val) {
    // Perform resource initialization, if any.
    // Register for attribute change notification, if needed.
}

void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPOpen, res_open);
    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}
```

Syntax for C++ close

```
void res_close(const char *res_name, void **attr_val);
```

You may select any name for the function.

Set the VCSAgValidateAndSetEntryPoint() parameter to the name of the entry point's function.

In the following example, the function res_close is defined as the Close entry point.

For example:

```
#include "VCSAgApi.h"

void res_close(const char *res_name, void **attr_val) {
    // Resource-specific de-initialization, if needed.
}
```

```
        // Unregister for attribute change notification, if any.
    }

void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPClose, res_close);
    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}
```

Syntax for C++ shutdown

```
void res_shutdown();
```

You may select any name for the function.

Set the `VCSAgValidateAndSetEntryPoint()` parameter to the name of the entry point's function.

In the following example, the function `res_shutdown` is defined as the Shutdown entry point.

For example:

```
#include "VCSAgApi.h"

void res_shutdown() {
    // Agent-specific de-initialization, if any.
}

void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPShutdown,
res_shutdown);
    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}
```

Agent framework primitives

Primitives are C++ methods implemented by the agent framework. The following sections define the primitives.

See also:

[“Agent Framework primitives for container support”](#) on page 82

[“Agent framework primitives for Windows”](#) on page 85

VCSAgRegisterEPStruct

Deprecated.

Use VCSAgValidateAndSetEntryPoint and VCSAgInitEntryPointStruct.

VCSAgSetCookie

Deprecated. See [“VCSAgSetCookie2”](#) on page 65.

VCSAgSetCookie2

```
void *VCSAgSetCookie2(const char *name, void *cookie)
```

This primitive requests the agent framework to store a cookie given by the void *cookie parameter. If there is a value already associated with the cookie, the primitive sets the new value and atomically returns the old value. If there is no value associated with the cookie, the primitive returns NULL.

This value, which is transparent to the agent framework, can be obtained by calling the primitive VCSAgGetCookie(). A cookie is not stored permanently. It is lost when the agent process exits. This primitive can be called from any entry point. For example:

```
#include "VCSAgApi.h"
...
// Assume that the online, offline, and monitor
// operations on resource require a certain key. Also
// assume that obtaining this key is time consuming, but
// that it can be reused until this process is
// terminated.
//
// In this example, the open entry point obtains the key
// and stores it as a cookie. Subsequent online,
// offline, and monitor entry points get the cookie and
// use the key.
//
// Note that the cookie name can be any unique string.
// This example uses the resource name as the cookie
// name.
//
```

```
void *get_key() {
    ...
}
void res_open(const char *res_name, void **attr_val) {
    if (VCSAgGetCookie(res_name) == NULL) {
        void *key = get_key();
        VCSAgSetCookie2(res_name, key);
    }
}
VCSAgResState res_monitor(const char *res_name, void
**attr_val, int *conf_level_ptr) {
    VCSAgResState state = VCSAgResUnknown;
    *conf_level_ptr = 0;
    void *key = VCSAgGetCookie(res_name);
    if (key == NULL) {
        // Take care of the rare cases when
        // the open entry point failed to
        // obtain the key and set the the cookie.
        key = get_key();
        VCSAgSetCookie2(res_name, key);
    }
    // Use the key for testing if the resource is
    // online, and set the state accordingly.
    ...
    return state;
}
```

VCSAgRegister

```
void  
VCSAgRegister(const char *notify_res_name,  
              const char *res_name,  
              const char *attr_name);
```

This primitive requests that the agent framework notify the resource `notify_res_name` when the value of the attribute `attr_name` of the resource `res_name` is modified. The notification is made by calling the `attr_changed` entry point for `notify_res_name`.

Note that `notify_res_name` must be the same as `res_name`.

This primitive can be called from any entry point, but it is useful only when the `attr_changed` entry point is implemented. For example:

```
#include "VCSAgApi.h"  
...  
void res_open(const char *res_name, void **attr_val) {  
  
    // Register to get notified when the  
    // "CriticalAttr" of this resource is modified.  
    VCSAgRegister(res_name, res_name, "CriticalAttr");  
  
    // Register to get notified when the "CriticalAttr"  
    // of current resource is modified. It is assumed  
    // that the name of the current resource is given  
    // as the first ArgList attribute.  
    VCSAgRegister((const char *) attr_val[0], (const  
                char *) attr_val[0], "CriticalAttr");  
  
}
```

VCSAgUnregister

```
void
VCSAgUnregister(const char *notify_res_name, const char
               *res_name,
               const char *attr_name);
```

This primitive requests that the agent framework stop notifying the resource `notify_res_name` when the value of the attribute `attr_name` of the resource `res_name` is modified. This primitive can be called from any entry point. For example:

```
#include "VCSAgApi.h"
...
void res_close(const char *res_name, void **attr_val) {

    // Unregister for the "CriticalAttr" of this resource.
    VCSAgUnregister(res_name, res_name, "CriticalAttr");

    // Unregister for the "CriticalAttr" of another
    // resource. It is assumed that the name of the
    // other resource is given as the first ArgList
    // attribute.
    VCSAgUnregister(res_name, (const char *)
        attr_val[0], "CriticalAttr");
}
```

VCSAgGetCookie

```
void *VCSAgGetCookie(const char *name);
```

This primitive requests that the agent framework get the cookie set by an earlier call to `VCSAgSetCookie2()`. It returns `NULL` if cookie was not previously set. This primitive can be called from any entry point. For example:

```
#include "VCSAgApi.h"
...
// Assume that the online, offline, and monitor
// operations on resource require a certain key. Also
// assume that obtaining this key is time consuming, but
// that it can be reused until this process is terminated.
//
// In this example, the open entry point obtains the key
// and stores it as a cookie. Subsequent online,
// offline, and monitor entry points get the cookie and
// use the key.
//
// Note that the cookie name can be any unique string.
// This example uses the resource name as the cookie name.
//

void *get_key() {
    ...
}

void res_open(const char *res_name, void **attr_val) {
    if (VCSAgGetCookie(res_name) == NULL) {
        void *key = get_key();
        VCSAgSetCookie2(res_name, key);
    }
}

VCSAgResState res_monitor(const char *res_name, void
    **attr_val, int *conf_level_ptr) {
    VCSAgResState state = VCSAgResUnknown;
    *conf_level_ptr = 0;
    void *key = VCSAgGetCookie(res_name);
    if (key == NULL) {
        // Take care of the rare cases when the open
        // entry point failed to obtain the key and
        // set the the cookie.
        key = get_key();
        VCSAgSetCookie2(res_name, key);
    }
    // Use the key for testing if the resource is
    // online, and set the state accordingly.
    ...
    return state;
}
```

VCSAgStrlcpy

```
void VCSAgStrlcpy(CHAR *dst, const CHAR *src, int size)
```

This primitive copies the contents from the input buffer “src” to the output buffer “dst” up to a maximum of “size” number of characters. Here, “size” refers to the size of the output buffer “dst.” This helps prevent any buffer overflow errors. The output contained in the buffer “dst” may be truncated if the buffer is not big enough.

VCSAgStrlcat

```
void VCSAgStrlcat(CHAR *dst, const CHAR *src, int size)
```

This primitive concatenates the contents of the input buffer “src” to the contents of the output buffer “dst” up to a maximum such that the total number of characters in the buffer “dst” do not exceed the value of “size.” Here, “size” refers to the size of the output buffer “dst.”

This helps prevent any buffer overflow errors. The output contained in the buffer “dst” may be truncated if the buffer is not big enough.

VCSAgSnprintf

```
int VCSAgSnprintf(CHAR *dst, int size, const char *format, ...)
```

This primitive accepts a variable number of arguments and works like the C library function “sprintf.” The difference is that this primitive takes in, as an argument, the size of the output buffer “dst.” The primitive stores only a maximum of “size” number of characters in the output buffer “dst.” This helps prevent buffer overflow errors. The output contained in the buffer “dst” may be truncated if the buffer is not big enough.

VCSAgCloseFile

```
void VCSAgCloseFile(void *vp)
```

Thread cleanup handler to close a file. The input (that is, vp) must be a file descriptor.

VCSAgDelString

```
void VCSAgDelString(void *vp)
```

Thread cleanup handler to delete a (char *). The input (vp) must be a pointer to memory allocated using “new char[xx]”.

VCSAgExec

```
int VCSAgExec(const char *path, char *const argv[], char *buf, long
buf_size, unsigned long *exit_codep)
```

Fork a new process, exec a program, wait for it to complete, and return the status. Also, capture the messages from stdout and stderr to buf. Caller must ensure that buf is of size $\geq$ buf_size.

VCSAgExec is a forced cancellation point. Even if the C++ entry point that calls VCSAgExec disables cancellation before invoking this API, the thread can get cancelled inside VCSAgExec. Therefore, the entry point must make sure that it pushes appropriate cancellation cleanup handlers before calling VCSAgExec. The forced cancellation ensures that a service thread running a timed-out entry point does not keep running or waiting for the child process created by this API to exit, but instead honors a cancellation request when it receives one.

Explanation of arguments to the function:

path	Name of the program to be executed.
argv	Arguments to the program. argv[0] must be same as path. The last entry of argv must be NULL. (Same as execv syntax)
buf	Buffer to hold the messages from stdout or stderr. Caller must supply it. This function will not allocate. When this function returns, buf will be NULL-terminated.
bufsize	Size of buf. If the total size of the messages to stdout/stderr is more than bufsize, only the first (buf_size - 1) characters will be returned.
exit_codep	Pointer to a location where the exit code of the executed program will be stored. This value should be interpreted as described by wait() on Unix Get Exit Code Process() on Windows

Return value: VCSAgSuccess if the execution was successful.

Example:

```
//
// ...
//
char **args = new char* [3];
char buf[100];
unsigned int status;

args[0] = "/usr/bin/ls";
args[1] = "/tmp";
args[2] = NULL;
```

```
int result = VCSAgExec(args[0], args, buf, 100, &status);

if (result == VCSAgSuccess) {

    // Windows NT:
    printf("Exit code of %s is %d\n", args[0], status);

    // Unix:
    if (WIFEXITED(status)) {
        printf("Child process returned %d\n", WEXITSTATUS(status));
    }
    else {
        printf("Child process terminated abnormally(%x)\n", status);
    }

}
else {
    printf("Error executing %s\n", args[0]);
}
//
// ...
//
```

VCSAgExecWithTimeout

```
int VCSAgExecWithTimeout(const char *path, char *const argv[],
    unsigned int timeout, char *buf, long buf_size, unsigned long
    *exit_codep)
```

Fork a new process, exec a program, wait for it to complete, return the status. If the process does not complete within the timeout value, kill it. Also, capture the messages from stdout or stderr to buf. The caller must ensure that buf is of size $\geq$ buf_size. VCSAgExecWithTimeout is a forced cancellation point. Even if the C++ entry point that calls VCSAgExecWithTimeout disables cancellation before invoking this API, the thread can get cancelled inside VCSAgExecWithTimeout. So the entry point needs to make sure that it pushes appropriate cancellation cleanup handlers before calling VCSAgExecWithTimeout. The forced cancellation ensures that a service thread running a timed out entry point does not keep running or waiting for the child process created by this API to exit but instead honors a cancellation request when it receives one.

Note: This API is not available on Windows.

Explanation of arguments to the function:

path	Name of the program to be executed.
argv	Arguments to the program. argv[0] must be same as path. The last entry of argv must be NULL. (Same as execv syntax).
timeout	Number of seconds within which the process should complete its execution. If zero is specified, this API defaults to VCSAgExec(), meaning the timeout is to be ignored. If the timeout value specified exceeds the time left for the entry point itself to timeout, the maximum possible timeout value is automatically used by this API. For example, if the timeout value specified in the API is 40 seconds, but the entry point itself times out after the next 20 seconds, the agent internally sets the timeout value for this API to 20-3=17 seconds. The 3 seconds are a grace period between the timeout for the process created using this API and the entry point process timeout.
buf	Buffer to hold the messages from stdout/stderr. The caller must supply it. This function does not allocate. When this function returns, buf is NULL-terminated.
bufsize	Size of buf. If the total size of the messages to stdout/stderr is more than bufsize, only the first (buf_size - 1) characters is returned.
exit_codep	Pointer to a location where the exit code of the executed program is stored. This value should interpreted as described by wait() on Unix

Return value: VCSAgSuccess if the execution is successful.

VCSAgSendTrap

```
void VCSAgSendTrap(const CHAR *msg)
```

This API is used to send a notification through the notifier process. The input (that is, msg) is the notification message to be sent.

VCSAgLockFile

```
int VCSAgLockFile(const char *fname, VCSAgLockType ltype,
VCSAgBlockingType btype, VCSAgErrnoType *errp)
```

Get a read or write (that is, shared or exclusive) lock on the given file. Both blocking and non-blocking modes are supported. Returns 0 if the lock could be obtained, or returns VCSAgErrWouldBlock if non-blocking is requested and the lock is busy. Otherwise returns -1. Each thread is considered a distinct owner of locks.

Mt-safe; deferred cancel safe.

Warning: Do not do any operations on the file (ex, open, or close) within this process, except through the VCSAgReadLockFile(), VCSAgWriteLockFile(), and VCSAgUnlockFile() interfaces.

VCSAgInitEntryPointStruct

```
void VCSAgInitEntryPointStruct(VCSAgAgentVersion agent_version)
```

This primitive enables agents to initialize the agent framework and the entry point struct depending on the agent framework version passed to this API.

Examples:

```
VCSAgInitEntryPointStruct(V50);  
VCSAgInitEntryPointStruct(V51);
```

If the agent registers with a version greater than V40, the following entry points are supported:

- online
- offline
- monitor
- clean
- open
- close
- attr_changed
- shutdown
- info
- action

For information on available registration version numbers, check the VCSAgApiDefs.h header file.

VCSAgSetStackSize

```
void VCSAgSetStackSize(int i)
```

The agent framework sets the default stack size for threads in agents to 1MB. Use VCSAgStackSize to set the calling thread's stack size to the specified value.

Note: This API is not available on Windows.

VCSAgUnlockFile

```
int VCSAgUnlockFile(const char *fname, VCSAgErrnoType *errp)
```

Release read or write (i.e shared or exclusive) lock on the given file. Returns 0, if the lock could be released, or else returns -1.

Mt-safe; deferred cancel safe.

Mt-safe; deferred cancel safe.

Warning: Do not do any operations on the file (ex, open, or close) within this process, except through the VCSAgReadLockFile(), VCSAgWriteLockFile(), and VCSAgUnlockFile() interfaces.

VCSAgDisableCancellation

```
int VCSAgDisableCancellation(int *old_statep)
```

If successful, return 0 and set old_statep to the previous cancellation state.

Note: This API is not available on Windows.

VCSAgRestoreCancellation

```
int VCSAgRestoreCancellation(int desired_state)
```

If successful, return 0 and the desired_state is set as the current cancellation state.

Note: This API is not available on Windows.

VCSAgSetEntryPoint

This primitive is deprecated.

See “[VCSAgValidateAndSetEntryPoint](#)” on page 77

VCSAgValidateAndSetEntryPoint

```
void VCSAgValidateAndSetEntryPoint(VCSAgEntryPoint ep, f_ptr)
```

This primitive enables an agent developer to register any C++ entry point with the agent framework.

VCSAgEntryPoint is an enumerated data type defined in VCSAgApiDefs.h.

For example:

```
VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, my_monitor_func);
```

VCSAgSetLogCategory

```
void VCSAgSetLogCategory(int cat_id)
```

Sets the log category of the agent to the value specified in cat_id.

VCSAgGetProductName

```
const CHAR *VCSAgGetProductName()
```

An API for C++ entry points to be able to get the name of the product for logging purposes.

VCSAgIsMonitorLevelOne

```
VCSAgBool VCSAgIsMonitorLevelOne();
```

See “[VCSAgIsMonitorLevelTwo](#)” on page 77.

VCSAgIsMonitorLevelTwo

```
VCSAgBool VCSAgIsMonitorLevelTwo();
```

VCSAgIsMonitorLevelOne and VCSAgIsMonitorLevelTwo can be used by the monitor entry point to find out whether the actions to be taken by the monitor entry point should correspond to basic monitoring or detailed monitoring.

These APIs work together with the LevelTwoMonitorFrequency attribute.

See “[Leveltwomonitorfrequency](#)” on page 170.

For example, if you set LevelTwoMonitorFrequency to 5, then every MonitorInterval seconds, VCSAgIsMonitorLevelOne() returns TRUE.

Additionally, every 5th monitor cycle, VCSAgIsMonitorLevelTwo() API also returns TRUE. This helps users configure detailed monitoring more effectively

and the agent developer does not need to make any calculations to find out when to execute the logic related to detail monitoring.

VCSAgMonitorReturn

```
VCSAgResState VCSAgMonitorReturn(VCSAgResState state, s32
conf_level, const CHAR *conf_msg)
```

VCSAgResState state: The state of the resource as found by the monitor entry point.

int conf_level: The confidence level with which the resource was found to be online. This can be a number from 10 to 100.

const char * conf_msg: If the resource is being reported as ONLINE from the monitor entry point with a confidence level lower than 100, this parameter accepts a string containing the reason for the lower confidence level for the resource. If the confidence level reported is 100 or if the resource state is reported as Offline or IntentionalOffline, the confidence message will get automatically cleared even if agent developer provides a confidence message string to this API.

VCSAgSetResEPTimeout

```
void VCSAgSetResEPTimeout(s32 tmo)
```

This API allows an agent entry point to extend its timeout value dynamically from within the entry point's execution context. This might be required if a command executed from the entry point takes longer than expected to complete and the entry point does not want to timeout. Symantec recommends using this API with caution because the intent of timeouts is to make sure that entry points finish on time.

VCSAgDecryptKey

```
VCSAgDecryptKey(char *key, char *outbuf, int buflen);
```

This API lets you decrypt an encrypted string passed in the ArgListValues by the user. Typically users encrypt string attribute values for passwords using the encryption commands provided by VCS One. An entry point can use this API to decrypt the encrypted string and get the original string.

VCSAgCloseFile

```
voidout VCSAgCloseFile(void *vp);
```

Use this API to push a thread cancellation cleanup handler for an open file descriptor. The API accepts a FILED file descriptor as parameter. The API closes the file descriptor passed as argument. It does not check if the fd is valid (≥ 0).

```
FILED fd = open("<filename>", ...);
if (fd >= 0) {
    pthread_cleanup_push(VCSAgCloseFile, (void *) fd);
    // Use fd
    pthread_cleanup_pop(1); // Run VCSAgCloseFile to close the file
    descriptor
}
else {
    // error, could not successfully open file
}
```

VCSAgCloseFilePtr

```
void VCSAgCloseFilePtr(void *vp)
```

Use this API to push a thread cancellation cleanup handler for an open file descriptor. The API accepts the address of a FILED file descriptor as parameter. The API closes the file descriptor whose address is passed as argument, only if the value of the descriptor is greater than or equal to 0.

```
FILED fd = -1;
pthread_cleanup_push(VCSAgCloseFilePtr, (void *) &fd);
FILED fd = open("<filename>", ...);
if (fd >= 0) {
    // Use fd
}
else {
    // error, could not successfully open file
}
pthread_cleanup_pop(1); // Run VCSAgCloseFilePtr to close the file
descriptor
```

VCSAgCloseFp

```
void VCSAgCloseFp(void *vp);
```

Use this API to push a thread cancellation cleanup handler for an open FILE*, usually returned by library calls like fopen() on UNIX. The API accepts a FILE* file pointer as parameter. The API closes the FILE*. It does not check if the FILE* is valid (non-NULL).

```
FILE * fp = fopen("<filename>", ...);
```

```
if (fp != NULL) {
    pthread_cleanup_push(VCSAgCloseFp, (void *) fp);
    // Use fp
    pthread_cleanup_pop(1); // Run VCSAgCloseFp to close the FILE *
}
else {
    // error, could not successfully open file
}
```

VCSAgCloseFpPtr

```
void VCSAgCloseFpPtr(void *vp)
```

Use this API to push a thread cancellation cleanup handler for an open FILE *, usually returned by library calls like `fopen()` on UNIX. The API accepts the address of a FILE * as parameter. The API closes the FILE * whose address is passed as argument, only if the value of the FILE * is non-NULL.

```
FILE * fp = NULL;
pthread_cleanup_push(VCSAgCloseFpPtr, (void *) &fp);
fp = fopen("<filename>", ...);
if (fp != NULL) {
    // Use fp
}
else {
    // error, could not successfully open file
}
pthread_cleanup_pop(1); // Run VCSAgCloseFpPtr to close the FILE *
```

VCSAgCloseDirp

```
void VCSAgCloseDirp(void *vp);
```

Use this API to push a thread cancellation cleanup handler for an open DIR *, usually returned by library calls like `fopen()` on UNIX. The API accepts a DIR * file pointer as parameter. The API closes the DIR *. It does not check if the DIR * is valid (non-NULL).

VCSAgCloseDirpPtr

```
void VCSAgCloseDirpPtr(void *vp);
```

Use this API to push a thread cancellation cleanup handler for an open DIR *, usually returned by library calls like `fopen()` on UNIX. The API accepts the address of a DIR * as parameter. The API closes the DIR * whose address is passed as argument, only if the value of the DIR * is non-NULL.

VCSAgDelPasswdPtr

```
void VCSAgDelPasswdPtr(void *vp);
```

This is a thread cancellation cleanup handler to delete a passwd *. The input must be a passwd **.

```
passwd *pp = NULL;
pthread_cleanup_push(VCSAgDelPasswdPtr, &pp);
pp = new passwd;
```

VCSAgDelString

```
void VCSAgDelString(void *vp);
```

Thread cleanup handler to delete a (char *). The input (i.e vp) must be a char* memory allocated using "new char[xx]".

```
char *cp = new char[10];
pthread_cleanup_push(VCSAgDelString, cp);
```

VCSAgDelStringPtr

```
void VCSAgDelStringPtr(void *vp);
```

Thread cleanup handler to delete a (char *). The input (i.e vp) must be a char** memory allocated using "new char[xx]".

```
char *cp = NULL;
pthread_cleanup_push(VCSAgDelStringPtr, &cp);
cp = new char[10];
```

This is the value of VCSONE_CONF environment variable if set or the default value if not set.

Agent Framework primitives for container support

The following APIs are for use in agents that run in AIX WPARs, XRM containers and Solaris zones. Note that zones are supported by Solaris version 10 and above.

Note that the earlier APIs for zones support are deprecated:

- VCSAgGetContainerName
- VCSAgGetContainerID
- VCSAgExecInContainer
- VCSAgISZoneCapable

VCSAgISContainerCapable

```
VCSAgBool VCSAgIsContainerCapable();
```

This API returns either True or False.

- For solaris zones
If the agent is running on a Solaris 10 (or higher version) system, the API returns True; otherwise it returns False.
Agents can use this API to decide whether or not to perform zone-specific operations like comparing the zone_id field in the psinfo structure with the ID of the zone name specified in the resource configuration to confirm whether the found process is indeed the process the agent is looking for.
- For XRM
If the agent is running on a system that has xrm available, the API returns True; otherwise it returns False.
- For WPARs
If the agent is running on a system that has WPARs available, the API returns True; otherwise it returns False.

VCSAgExecInContainerWithTmo

This API is deprecated.

See “[VCSAgExecInContainerWithTimeout](#)” on page 82.

VCSAgExecInContainerWithTimeout

```
int VCSAgExecInContainerWithTimeout((const CHAR *path, CHAR *const  
argv[] u32 timeout, CHAR *buf, long buf_size, unsigned long  
*exit_codep);
```

This API is similar to the VCSAgExecWithTimeout API. This API can be used by an agent only to execute a particular command or script in a specific container on the system. If there are no containers configured on the system, or if the agent has no need to exec a script in a specific container, the VCSAgExecWithTimeout API should be used.

Memory for buf and exit_codep should be allocated by the calling function.

VCSAgGetUID

```
int VCSAgGetUID(const CHAR *user, int *uid, int *euid, int
*home_exists);
```

This API checks if the given user is valid inside the container as specified in the resource object. The API returns the uid and euid of the user either inside the container if container info is set for the resource or on the global container if container info is not set for the resource. The home_exists parameter indicates if the specified user's home directory exists within the container.

Memory for uid, euid and home_exists must be allocated by the calling function. The API returns 0 on success and 1 on failure.

VCSAgIsPidInContainer

```
int VCSAgIsPidInContainer(VCSPID pid);
```

This API checks if the given pid is running inside the container as specified in the resource object.

Return values

- 1 if the proc pid is running inside the container
- 0 if the proc pid is not running inside the container
- -1 if the API cannot verify the container info for the process. This is possible if ContainerType is an invalid value.

VCSAgIsProcInContainer

```
int VCSAgIsProcInContainer(void *psinfo);
```

This API checks if the process corresponding to the given psinfo structure is running inside the container as specified in the resource object.

Return values

- 1 if the proc pid is running inside the container
- 0 if the proc pid is not running inside the container

- -1 if the API cannot verify the container info for the process. This is possible if ContainerType is an invalid value.

VCSAgGetContainerID2

```
int VCSAgGetContainerID2()
```

This API retrieves the ID of the container.

Based on the thread that is implementing the entry point, the agent identifies the resource for which this API is invoked and returns the container ID for that resource. The container ID is the ID of the container specified in the ContainerInfo attribute as the value of the Name key.

Return Values

- -1, if the resource or container name is NULL or the container is DOWN or the container is not applicable to the OS version the agent is running on.
- Non-negative container-id, if the container name is valid and the container is UP.

VCSAgGetContainerName2

```
char *VCSAgGetContainerName2();
```

For Solaris Zones, this API retrieves the name of the container, if set for the specified resource.

For XRM, the API retrieves the name of the Execution Context.

For WPARs, the API retrieves the name of the WPAR.

The API returns a pointer to the container name. It is the responsibility of the caller to free the memory associated with the returned pointer.

The name of the container is the value set in the group-level attribute ContainerInfo for the group the resource belongs to.

Agent framework primitives for Windows

This chapter describes agent primitives for use by agents in a Windows environment. These primitives are used for:

- Logging messages
- Converting ArgList values from UTF-8 encoding to UCS-2 encoding (wchar)
- Manipulating UCS-2 (wchar) arguments
- Allocating and freeing memory

C++ logging primitives for use by Windows agents

The four C++ logging primitives provided by the agent framework include the two application logging macros:

```
VCSAG_CONSOLE_LOG_MSG (sev,msgid,flags,fmt,variable_args ... )
VCSAG_LOG_MSG (sev,msgid,flags,fmt,variable_args ... )
```

and the two debugging macros:

```
VCSAG_LOGDBG_MSG (dbgsev,flags,fmt,variable_args ... )
VCSAG_RES_LOG_MSG (dbgsev,flags,fmt,variable_args ... )
```

See “[C++ agent logging APIs](#)” on page 108.

Using logging primitives with localized Windows agents

The C++ logging primitives can be used to log Windows agent messages that use wchar arguments. The agents can simply pass in a format string in wchar as well as arguments in wchar to the log macros. The agent framework converts the wchar message to UTF-8 encoding before writing them to log files. The log macros accept a variable number of arguments as input. The arguments can be wchar string, wchar, integer, or floating point values.

Defining UCS-2 encoding when compiling agents

To enable the logging primitives to log messages using UCS-2 encoding (wchar), compile the agent with the “UNICODE” flag defined. Otherwise, compilation fails with errors that indicate a mismatch between the expected argument types and the actual UCS-2 encoded messages sent as arguments.

A sample rule to use to build an agent named MyAgent

```
.
.
MyAgent.o:    MyAgent.C
              $(CCPLUS) /FoMyAgent.o -DUNICODE -D_UNICODE -c MyAgent.C
.
.
MyAgent:      MyAgent.o
```

```
$(CCLINK) /dll /OUT:MyAgent MyAgent.o $(NEED_LIBS)
```

where CCPLUS, CCLINK, NEED_LIBS are set as required for the compilation and linking of the agent code.

C++ primitives for ArgList conversion

The two primitives provided for converting the UCS-2 encoded `ArgList` attribute are:

- `VCSAgGetEncodedArgList`
- `VCSAgDelEncodedArgList`.

VCS One on Windows allows users to use localized values for attributes. Therefore, Windows agents must handle UCS-2 encoded or OS encoded values. All arguments in the `ArgList` attribute are passed to the Windows agents in UTF-8 encoding, which the agent must convert to UCS-2 encoding or OS encoding. The `VCSAgGetEncodedArgList` is provided for this purpose.

`VCSAgGetEncodedArgList` allocates memory as needed to the encoded `ArgList` array output. This memory can be freed by the caller using the primitive `VCSAgDelEncodedArgList`.

See “[VCSAgDelEncodedArgList](#)” on page 87).

VCSAgGetEncodedArgList

```
int VCSAgGetEncodedArgList (VCSAgEncodingType from_encoding,
void
    **input_arglist_values, VCSAgEncodingType to_encoding, void
    ***output_arglist_values)
```

This primitive converts the `ArgList` values in UTF-8 encoding to UCS-2 or the OS encoding. `VCSAgGetEncodedArgList` stores a pointer to the encoded array in the “`output_arglist_values`” and returns 0 on success and 1 on failure.

The input `ArgList` array must be terminated with a NULL string to indicate the end of the input `ArgList`.

Note: This primitive only supports conversions to UCS-2 and OS encoding.

VCSAgDelEncodedArgList

```
void VCSAgDelEncodedArgList (void **enc_arg_list)
```

This primitive must be used to free the memory allocated by the agent in the primitive `VCSAgGetEncodedArgList`. This primitive expects the input, `enc_arg_list`, to be terminated by a NULL string.

`VCSAgDelEncodedArgList` can be called from any entry point that is passed the `ArgList` values.

C++ primitives for ArgList string manipulation

Three primitives provided for manipulating `wchar` strings in UCS-2 encoding are:

- `VCSAgStrncpyW`
- `VCSAgStrlcatW`
- `VCSAgSnprintfW`

VCSAgStrncpyW

```
void VCSAgStrncpyW (wchar *dst, const wchar *src, int size)
```

This primitive must be used by agents to perform string copy operations on `wchar` arguments. The argument “size” refers to the number of characters in the destination buffer “dst.” This helps prevent buffer overflow errors. The output contained in the buffer “dst” may be truncated if the buffer is not large enough.

VCSAgStrlcatW

```
void VCSAgStrlcatW (wchar *dst, const wchar *src, int size)
```

This primitive must be used in the agents to perform a string concatenation operation of `wchar` arguments. The argument “size” refers to the number of characters in the destination buffer “dst.” This helps prevent buffer overflow errors. The output contained in the buffer “dst” may be truncated if the buffer is not large enough.

VCSAgSnprintfW

```
int VCSAgSnprintfW (wchar *buffer, int size, const wchar
*format, ... )
```

This primitive accepts a variable number of arguments and works like the C library function “`vswprintf`,” except that `VCSAgSnprintfW` takes in the size of the output buffer “dst” as an argument. The primitive stores only a maximum of “size” number of characters in the output buffer “dst.” This

helps prevent buffer overflow errors. The output contained in the buffer “dst” may be truncated in the buffer is not large enough.

C++ primitives for allocating and freeing memory

Two primitives are provided for use by Windows agents to allocate and free memory:

- VCSAgAllocMem
- VCSAgDeleteMem

VCSAgAllocMem

```
void *VCSAgAllocMem (size_t size)
```

This primitive can be used to allocate “size” bytes of memory. The primitive VCSAgDeleteMem must be used to free the memory that is allocated using this API.

VCSAgDeleteMem

```
void VCSAgDeleteMem (void *mp)
```

This primitive can be used to free memory allocated to “mp.” The memory for “mp” should have been allocated in a previous call to the API VCSAgAllocMem.

Example: Windows-based primitives for the FileOnOff agent

The following example illustrates the use of the Windows-based primitives within an agent for FileOnOff.

```
/* @(#)src/agent/w2k/FileOnOff/FileOnOff.C      2.4 01/04/03
11:57:11 - */
static const char sccsid[] =
"@(#)VCS:src/agent/w2k/FileOnOff/FileOnOff.C 2.4";
/*
 * START COPYRIGHT-NOTICE: 1998, 2003
 *
 * Copyright 2008 Symantec Corporation.
 * All rights reserved.
 *
 */
#include <stdio.h>
#include <windows.h>
#include "VCSAgApi.h"

WCHAR *ConvertToUnicodeString(const char *attr_val)
{
    int len = 0,outLen ;
    WCHAR * szOutStr = NULL;
```

```

        int num_bytes = 0, num_written = 0;
        len = strlen(attr_val);

        // Call the API to convert the UTF-8 encoded string into
a
        // UCS-2 encoded string
        VCSAgEncodeString(VCSAgUTF8, /* convert from UTF8 */
                          (void *)attr_val,
                          len,
                          VCSAgUCS2, /* to the UCS2 format */
                          (void **)&szOutStr,
                          &outLen);

        return szOutStr;
    }
    //////////////////////////////////////
    //////////////////////////////////////
    // Function:      res_monitor
    //
    // Description: Determine if the given file is online (file
    exists)
    //
    //                or offline (file does not exist).
    //
    // Notes:         None
    //////////////////////////////////////
    //////////////////////////////////////

VCSAgResState res_monitor(const char *res_name, void **attr_val,
int *conf_level)
{
    VCSAG_LOG_INIT("res_monitor");

    VCSAgResState state;
    DWORD attrs ;

    // Temporary variables for sample code
    WCHAR * szPath;
    WCHAR *mymem = NULL;
    WCHAR **mymem2 = NULL;
    WCHAR **new_args = NULL;
    int ret = 0;
    int indx = 0;
    int temp_number = 0;

    // Dummy structure defined for the purpose of example
code
    struct mystruct{
        int a;
        WCHAR b[32];
    } *inst1;

```

```

/*****
***
    * Sample code for using the new AGFW APIs
*****/

// Allocate 32 bytes of type WCHAR for mymem
mymem = VCSAgNew(WCHAR, 32);

if (mymem) {
    // Memory allocation succeeded
    VCSAgStrncpyW(mymem, I, "copy ", 32);

    VCSAgStrncatW(mymem, I, "concatenate", 32);

    VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
        I, "Sample string is %s", mymem);

    temp_number = 10;
    VCSAgSnprintfW(mymem, 32,
        I, "Testing VCSAgSnprintfW...temp_number is
        %d\n", temp_number);

    VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
        I, "Sample string is %s", mymem);

    // Delete memory allocated to mymem
    VCSAgDelete((void *)mymem);
    mymem = NULL;
} else {
    // Memory allocation failed
    VCSAG_LOG_MSG(VCS_NOTICE, 1005,
VCS_DEFAULT_FLAGS,
        I, "Failed to allocate memory for mymem");
}

// Allocating memory for a struct now

inst1 = VCSAgNew(struct mystruct, 1);

if (inst1) {
    // Memory allocation succeeded
    VCSAgStrncpyW(inst1->b,
        I, "Copying dummy string", 32);

    VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
        I, "Sample string is %s", inst1->b);

    // Delete memory allocated to inst1

```

```

        VCSAgDelete(inst1);

    } else {
        // Memory allocation failed
        VCSAG_LOG_MSG(VCS_NOTICE, 1006,
VCS_DEFAULT_FLAGS,
                        I,"Failed to allocate memory for inst1");
    }

    // Allocating memory for an array of pointers
    // Here, there will be four entries in the array
    myptr = VCSAgNew(WCHAR *, 4);
    if (myptr) {
        // Memory allocation succeeded
        myptr[0] = VCSAgNew(WCHAR, 16);
        VCSAgStrlcpyW(myptr[0], I,"One", 16);

        myptr[1] = VCSAgNew(WCHAR, 16);
        VCSAgStrlcpyW(myptr[1], I,"Two", 16);

        VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
                        I,"Contents of myptr: %s and %s",
                        myptr[0], myptr[1]);

        // Free the memory allocated
        VCSAgDelete(myptr[0]);
        VCSAgDelete(myptr[1]);
        VCSAgDelete(myptr);
    }

    // Call the function that uses the VCSAgEncodeString API
    szPath = ConvertToUnicodeString((char *)attr_val[0]);
    VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
        I,"Encoded string is %s", szPath);

    // Sample usage of VCSAgGetEncodedArgList API
    ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val,
        VCSAgUCS2, (void ***)&new_args);

    if (ret) {
        // Encode API failed
        VCSAG_LOG_MSG(VCS_NOTICE, 1007,
VCS_DEFAULT_FLAGS,
                        I,"Failed to encode the ArgList entries");
    } else {
        // Encode succeeded
        VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
            I,"The encoded arg list values are:");
        while (new_args[indx] != NULL) {
            VCSAG_LOGDBG_MSG(VCS_DBG1,
VCS_DEFAULT_FLAGS,
                        I,"arg %d is %s", indx, new_args[indx]);

```

```

        indx++;
    }

    // Sample usage of API to delete the encoded
    // ArgList entries

    VCSAgDelEncodedArgList((void **)new_args);

}

/*****
***
    * End of sample code for using the new AGFW APIs
*****/

:
// Actual code here
:
:

VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
    I, "Monitor is successful");

if(szPath)
    VCSAgDelString((void *)szPath);

return state;
}

:
// Rest of the agent code
:
:
```

Creating entry points in scripts

- [About creating entry points in scripts](#)
- [Syntax for script entry points](#)
- [Example script entry points](#)

About creating entry points in scripts

On UNIX, script agents use the Script51Agent binary that are shipped with the product. The Script51Agent binaries are located at:

`$VCSONE_HOME/bin/Script51Agent`

On Windows you may use the Default50Agent.dll.

`%VCSONE_HOME%\bin\Default50Agent.dll`

You must implement the VCSAgStartup entry point using C++.

You may implement other entry points using C++ or scripts. If you implement no other entry points in C++, the VCSAgStartup entry point is not required.

On UNIX, use the Script51Agent binary to develop all entry points as scripts.

You can use Perl or shell scripts to develop entry points.

On Windows, use VCSdefault.dll.

See “[Implementing entry points using scripts](#)” on page 145.

Rules for using script entry points

Script entry points can be executables or scripts, such as shell or Perl (the product includes a Perl distribution).

Adhere to the following rules when implementing a script entry point:

On the Windows platform

- The name of the script file must be the same as the entry point, followed by a period and the appropriate extension. The extension assumes the value `exe` for executable programs, `sh` for shell scripts compatible with MKS Toolkit, `bat` or `cmd` for batch files, `pl` for Perl scripts, `js` for Javascript, or `vbs` for Visual Basic scripts.
- Place the file in the directory `%VCSONE_HOME%\bin\resource_type`. For example, if the `online` entry point for Oracle is implemented using Perl, the online script must be:

`%VCSONE_HOME%\bin\Oracle\online.pl`

On UNIX platforms

- In the VCSAgStartup entry point, if you do not set a C++ function for an entry point using the VCSAgValidateAndSetEntryPoint() API, then the agent framework assumes the entry point is script-based. See “[About the VCSAgStartup routine](#)” on page 37.
- Verify the name of the script file is the same as the entry point .

- Place the file in the directory `$VCSONE_HOME/bin/resource_type`. If, for example, the `online` script for Oracle were implemented using Perl, the `online` script must be:


```
$VCSONE_HOME/bin/Oracle/online.pl
```
- If you write scripts in shell, verify the `PATH` environment variable includes the directory where `sh` is installed.

Parameters and values for script entry points

The input parameters of script entry points are passed as command-line arguments. The first command-line argument for all the entry points is the name of the resource (except `shutdown`, which has no arguments).

Some entry points have an output parameter that is returned through the program exit value. See the entry point description for more information.

See [“Syntax for script entry points”](#) on page 97.

ArgList attributes

See [“About the ArgList and ArgListValues attributes”](#) on page 39.

Example

If the type definition of type “Foo” is:

```
<agent name="Foo" version="1.0" bundled="False">
  <attributes>
    <IntAttr type="int" dimension="scalar">
      .
    </IntAttr>
    <StringAttr type="str" dimension="scalar">
      .
    </StringAttr>
    <VectorAttr type="str" dimension="vector">
      .
    </VectorAttr>
    <AssocAttr type="str" dimension="assoc">
      .
    </AssocAttr>
  </attributes>
  <attributes class=static>
    <ArgList type="str" dimension="vector">
      <default>
        <val>IntAttr</val>
        <val>StringAttr</val>
        <val>VectorAttr</val>
        <val>AssocAttr</val>
      </default>
    </ArgList>
```

```
    </attributes>  
</agent>
```

And the resource is defined in the configuration file `main.xml` as:

```
<resources>  
  <resource name="Bar" type="Foo">  
    <attribute IntAttr>  
      <scalar>100</scalar>  
    </attribute>  
    <attribute StringAttr>  
      <scalar>"Oracle"</scalar></attribute>  
    <attribute VectorAttr>  
      <val>"vol1"</val>  
      <val>"vol2"</val>  
    </attribute>  
    <attribute AssocAttr>  
      <val key="disk1">1024</val>  
      <val key="disk2">512</val>  
    </attribute>  
  </resource>  
</resources>
```

For Script51Agent, the online script for Bar is invoked as:

```
online Bar IntAttr 1 100 StringAttr 1 Oracle VectorAttr 3 vol1  
vol2 vol3 AssocAttr 4 disk1 1024 disk2 512
```

Syntax for script entry points

The following paragraphs describe the syntax for script entry points.

Syntax for the monitor script

```
monitor resource_name ArgList_attribute_values
```

A script entry point combines the status and the confidence level in the exit value. For example:

- 99 indicates unknown.
- 100 indicates offline.
- 101 indicates online and a confidence level of 10.
- 102–109 indicates online and confidence levels 20–90.
- 110 indicates online and confidence level 100.
- 200 indicates intentional offline.

If the exit value is not one of the above values, the status is considered unknown.

Syntax for the online script

```
online resource_name ArgList_attribute_values
```

The exit value is interpreted as the expected time (in seconds) for the online procedure to be effective. It also means the time (in seconds) that must pass before executing the monitor entry point to validate proper operation. The exit value is typically 0.

Syntax for the offline script

```
offline resource_name ArgList_attribute_values
```

The exit value is interpreted as the expected time (in seconds) for the offline procedure to be effective. The exit value is typically 0.

Syntax for the clean script

```
clean resource_name clean_reason argList_attribute_values
```

The variable *clean_reason* equals one of the following values:

- 0 - The *offline* entry point did not complete within the expected time.
(See “[OfflineTimeout](#)” on page 173.)
- 1 - The *offline* entry point was ineffective.
- 2 - The *online* entry point did not complete within the expected time.(See “[OnlineTimeout](#)” on page 174.)
- 3 - The *online* entry point was ineffective.
- 4 - The resource was taken offline unexpectedly.
- 5 - The *monitor* entry point consistently failed to complete within the expected time.
(See “[FaultOnMonitorTimeouts](#)” on page 169.)

The exit value is 0 (successful) or 1.

Note: Because entry point timeouts cannot be set for agents running on Windows, *clean_reason* values of 0, 2, and 5 do not apply.

Syntax for the action script

```
action resource_name  
ArgList_attribute_values_AND_action_arguments
```

The exit value is 0 (successful) or 1 (if unsuccessful).

The agent framework limits the action entry point output to 2048 bytes.

Syntax for the attr_changed script

```
attr_changed resource_name changed_resource_name  
changed_attribute_name new_attribute_value
```

The exit value is ignored.

Note: This entry point is called only if you register for change notification using the primitive `VCSAgRegister()` (see “[VCSAgRegister](#)” on page 67), or the agent parameter `RegList` (see “[RegList](#)” on page 175).

Syntax for the info script

```
info resource_name resinfo_op ArgList_attribute_values
```

The attribute *resinfo_op* can have the values 1 or 2.

Values of <i>resinfo_op</i>	Significance
1	Add and initialize static and dynamic name-value data pairs in the <code>ResourceInfo</code> attribute.
2	Update just the dynamic data in the <code>ResourceInfo</code> attribute.

This entry point can add and update static and dynamic name-value pairs to the `ResourceInfo` attribute. The `info` entry point has no specific output, but rather, it updates the `ResourceInfo` attribute.

Syntax for the open script

```
open resource_name ArgList_attribute_values
```

The exit value is ignored.

Syntax for the close script

```
close resource_name ArgList_attribute_values
```

The exit value is ignored.

Syntax for the shutdown script

```
shutdown
```

The exit value is ignored.

Example script entry points

The following example shows entry points written in a shell script.

Online entry point for FileOnOff

The FileOnOff example entry point is simple. When the agent's `online` entry point is called by the agent, the entry point expects the name of the resource as the first argument, followed by the values of the remaining `ArgList` attributes.

- For agents that are registered as less than V50, the entry point expects the values of the attributes in the order the attributes have been specified in the `ArgList` attribute.
- For agents registered as V50 and greater, the entry point expects the `ArgList` in tuple format: the name of the attribute, the number of elements in the attribute's value, and the value.

Note: The actual VCS One FileOnOff entry points are written in C++, but for this example, shell script is used.

Monitor entry point for FileOnOff

When the agent's `monitor` entry point is called by the agent, the entry point expects the name of the resource as the first argument, followed by the values of the remaining `ArgList` attributes.

- For agents that are registered as less than V50, the entry point expects the values of the attributes in the order the attributes have been specified in the `ArgList` attribute.
- For agents registered as V50 and greater, the entry point expects the `ArgList` in tuple format: the name of the attribute, the number of elements in the attribute's value, and the value.

If the file exists it returns exit code 110, indicating the resource is online with 100% confidence. If the file does not exist the monitor returns 100, indicating the resource is offline. If the state of the file cannot be determined, the monitor returns 99.

```
#!/bin/sh
# FileOnOff Monitor script
# Expects Resource Name and Pathname

. $VCSONE_HOME/bin/ag_i18n_inc.sh
RESNAME=$1
VCSAG_SET_ENVS $RESNAME
#check if second attribute provided
#Exit with unknown and log error if not provided.
if [ -z "$2" ]
then
    VCSAG_LOG_MSG "W" "The value for PathName is not specified"\
    exit 99
else
    if [ -f $2 ]; then exit 110;
    # Exit online (110) if file exists
    # Exit offline (100) if file does not exist
    else exit 100;
    fi
fi
```

Monitor entry point with intentional offline

This script includes the intentional offline functionality for the `MyCustomApp` agent.

See “[About on-off, on-only, and persistent resources](#)” on page 16.

Note that the method to detect intentional offline of an application depends on the type of application. The following example assumes that the application writes a status code into a file if the application is intentionally stopped.

```
#!/bin/sh

. ${CLUSTER_HOME}/bin/ag_i18n_inc.sh

ResName=$1; shift;
VCSAG_SET_ENVS $ResName
// Obtain the attribute values from ArgListValues
parse_arglist_values();
RETVAL=$?

if [ ${RETVAL} -eq ${VCSAG_RES_UNKNOWN} ]; then
    // Could not get all the required attributes from
    ArgListValues
    exit $VCSAG_RES_UNKNOWN;
fi

// Check if the application's process is present in the ps
// output
check_if_app_is_running();
RETVAL=$?

if [ ${REVAL} -eq ${VCSAG_RES_ONLINE} ]; then
    // Application process found
    exit $VCSAG_RES_ONLINE;
fi

// Application process was not found; Check if user gracefully
// shutdown the application
grep "MyCustomAppCode 123 : User initiated shutdown command"
${APPLICATION_CREATED_STATUS_FILE}
RETVAL=$?

if [ ${REVAL} -eq 0 ]; then
    // Found MyCustomAppCode 123 in the application's status
    // file that gets created by the application on graceful
    // shutdown
    exit $VCSAG_RES_INTENTIONALOFFLINE;
else
    // Did not find MyCustomAppCode 123; hence application has
    // crashed or gone down unintentionally
    exit $VCSAG_RES_OFFLINE;
fi

// Monitor should never come here
exit $VCSAG_RES_UNKNOWN;
```

Offline entry point for FileOnOff

When the agent's `offline` entry point is called by the agent, the entry point expects the name of the resource as the first argument, followed by the values of the remaining `ArgList` attributes.

- For agents that are registered as less than V50, the entry point expects the values of the attributes in the order the attributes have been specified in the `ArgList` attribute.
- For agents registered as V50 and greater, the entry point expects the `ArgList` in tuple format: the name of the attribute, the number of elements in the attribute's value, and the value.

When the values are accepted by the entry point, the file is deleted.

```
#!/bin/sh
# FileOnOff Offline script
# Expects ResourceName and Pathname
#
. $VCSONE_HOME/bin/ag_i18n_inc.sh
RESNAME=$1
VCSAG_SET_ENVS $RESNAME
#check if second attribute provided
if [ -z "$2" ]
then
    VCSAG_LOG_MSG "W" "The value for PathName is not specified"\
    1020
else
    #remove the file
    /bin/rm -f $2
fi
exit 0;
```


Logging agent messages

- [About logging agent messages](#)
- [Logging in C++ and script-based entry points](#)
- [C++ agent logging APIs](#)
- [Script entry point logging functions](#)

About logging agent messages

This chapter describes APIs and functions that developers can use within their custom agents to generate log file messages conforming to a standard message logging format.

See: [Appendix A, “Using VCS agents with VCS One”](#) on page 209

Logging in C++ and script-based entry points

Developers creating C++ agent entry points can use a set of macros for logging application messages or debug messages. Developers of script-based entry points can use a set of methods, or “wrappers,” that call the `halog` utility to generate application or debug messages.

Symantec recommends using the `ag_i18n_inc` subroutines for logging. The subroutines set the category ID for the messages and provide a header for the log message, which includes the resource name and the entry point name.

Agent messages: format

An agent log message consists of five fields. The format of the message is:

<Timestamp> <Mnemonic> <Severity> <UMI> <MessageText>

The following is an example message, of severity ERROR, generated by the FileOnOff agent's online entry point. The message is generated when the agent attempts to bring online a resource, a file named "MyFile":

```
Jun 26 2003 11:32:56 VCS One ERROR V-16-2001-14001
FileOnOff:MyFile:online:Resource could not be brought up
because,the attempt to create the file (filename) failed
with error (Is a Directory)
```

The first four fields of the message above consists of the *timestamp*, an uppercase *mnemonic* that represents the product, the *severity*, and the *UMI* (unique message ID). The subsequent lines contain the *message text*.

Timestamp

The timestamp indicates when the message was generated. It is formatted according to the locale.

The time is the time on the client where the log message was generated.

Mnemonic

The mnemonic field is used to indicate the product.

The mnemonic, must use all capital letters. All bundled agents, enterprise agents, and custom agents use the mnemonic: "VCS One"

Severity

The severity of each message displays in the third field of the message (Critical, Error, Warning, Notice, or Information for normal messages; 1-21 for debug messages). All C++ logging macros and script-based logging functions provide a means to define the severity of messages, both normal and debugging.

UMI

The UMI (unique message identifier) includes an originator ID, a category ID, and a message ID.

- The originator ID is a decimal number preceded by a "V-" that defines the product that the message comes from. This ID is assigned by Symantec.
- The category ID is a number in the range of 0 to 65536 assigned by Symantec. The category ID indicates the agent that message came from. For each custom agent, you must contact Symantec so that a unique category ID can be registered for the agent.

- For C++ messages, the category ID is defined in the `VCSAgStartup` entry point.
See “[Log category](#)” on page 113.
- For script-based entry points, the category is set within the `VCSAG_SET_ENVS` function
See “[VCSAG_SET_ENVS](#)” on page 118.
- For debug messages, the category ID, which is 50 by default, need not be defined within logging functions.
- Message IDs can range from 0 to 65536 for a category ID. Each normal message (that is, non-debug message) generated by an agent must be assigned a message ID. For C++ entry points, the `msgid` is set as part of the `VCSAG_LOG_MSG` and `VCSAG_CONSOLE_LOG_MSG` macros. For script-based entry points, the `msgid` is set using the `VCSAG_LOG_MSG` function. The `msgid` field is not used by debug functions or required in debug messages.
See “[VCSAG_LOG_MSG](#)” on page 121.

Message text

The message text is a formatted message string preceded by a dynamically generated header consisting of three colon-separated fields. namely, *<name of the agent>:<resource>:<name of the entry point>:<message>*. For example:

```
FileOnOff:MyFile:online:Resource could not be brought up
because,the attempt to create the file (MyFile) failed
with error (Is a Directory)
```

- In the case of C++ entry points, the header information is generated.
- In the case of script-based entry points, the header information is set within the `VCSAG_SET_ENVS` function (see “[VCSAG_SET_ENVS](#)” on page 118).

C++ agent logging APIs

The agent framework provides four logging APIs (macros) for use in agent entry points written in C++.

These APIs include two application logging macros:

```
VCSAG_CONSOLE_LOG_MSG(sev, msgid, flags, fmt, variable_args...)
VCSAG_LOG_MSG(sev, msgid, flags, fmt, variable_args...)
```

and the macros for debugging:

```
VCSAG_LOGDBG_MSG(dbgsev, flags, fmt, variable_args...)
VCSAG_RES_LOG_MSG(dbgsev, flags, fmt, variable args...)
```

Agent application logging macros for C++ entry points

You can use the macro `VCSAG_LOG_MSG` within C++ agent entry points to log all messages ranging in severity from `CRITICAL` to `INFORMATION` to the agent log file. Use the `VCSAG_CONSOLE_LOG_MSG` macro to send messages to the `vcsonclientd` log. Where the messages are of `CRITICAL` or `ERROR` severity, the message is also logged to the console of the client system.

The following table describes the argument fields for the application logging macros:

<code>sev</code>	Severity of the message from the application. The values of <code>sev</code> are macros <code>VCS_CRITICAL</code> , <code>VCS_ERROR</code> , <code>VCS_WARNING</code> , <code>VCS_NOTICE</code> , and <code>VCS_INFORMATION</code> ; see “Severity arguments for C++ macros” on page 111.
<code>msgid</code>	The 16-bit integer message ID.
<code>flags</code>	Default flags (0) prints UMI, NEWLINE. A macro, <code>VCS_DEFAULT_FLAGS</code> , represents the default value for the flags.
<code>fmt</code>	A formatted string containing formatting specifiers symbols. For example: “Resource could not be brought down because the attempt to remove the file (%s) failed with error (%d)”
<code>variable_args</code>	Variable number (as many as 6) of type <code>char</code> , <code>char *</code> , or integer

In the following example, the macros are used to log an error message to the agent log and to the console:

```
.
.
VCSAG_LOG_MSG(VCS_ERROR, 14002, VCS_DEFAULT_FLAGS,
    "Resource could not be brought down because the
    attempt to remove the file(%s) failed with error(%d)",
    (CHAR *) (*attr_val), errno);

VCSAG_CONSOLE_LOG_MSG(VCS_ERROR, 14002, VCS_DEFAULT_FLAGS,
    "Resource could not be brought down because, the
    attempt to remove the file(%s) failed with error(%d)",
    (CHAR *) (*attr_val), errno);
```

Agent debug logging macros for C++ entry points

Use the macros `VCSAG_RES_LOG_MSG` and `VCSAG_LOGDBG_MSG` within agent entry points to log debug messages of a specific severity level to the agent log.

Use the `LogDbg` attribute to specify a debug message severity level. See the description of the `LogDbg` attribute ([“LogDbg”](#) on page 170). Set the `LogDbg` attribute at the resource type level. The attribute can be overridden to be set at the level for a specific resource.

The `VCSAG_LOGDBG_MSG` macro controls logging at the level of the resource type level, whereas `VCSAG_RES_LOG_MSG` macro can enable logging debug messages at the level of a specific resource.

The following table describes the argument fields for the application logging macros:

<code>dbgsev</code>	Debug severity of the message. The values of <code>dbgsev</code> are macros ranging from <code>VCS_DBG1</code> to <code>VCS_DBG21</code> . See “Severity arguments for C++ macros” on page 111.
<code>flags</code>	Describes the logging options. Default flags (0) prints UMI, NEWLINE. A macro, <code>VCS_DEFAULT_FLAGS</code> , represents the default value for the flags
<code>fmt</code>	A formatted string containing symbols. For example: “PathName is (%s)”
<code>variable_args</code>	Variable number (as many as 6) of type <code>char</code> , <code>char *</code> or integer

For example:

```
VCSAG_RES_LOG_MSG(VCS_DBG4, VCS_DEFAULT_FLAGS, "PathName is (%s)",  
                 (CHAR *) (*attr_val));
```

For the example shown, the specified message is logged to the agent log if the specific resource has been enabled (that is, the `LogDbg` attribute is set) for logging of debug messages at the severity level `DBG4`.

Severity arguments for C++ macros

A severity argument for a logging macro, for example, `VCS_ERROR` or `VCS_DBG1`, is in fact a macro itself that expands to include the following information:

- actual message severity
- function name
- name of the file that includes the function
- line number where the logging macro is expanded

For example, the application severity argument `VCS_ERROR` within the `monitor` entry point for the `FileOnOff` agent would expand to include the following information:

```
ERROR, res_monitor, FileOnOff.C, 28
```

Application severity macros map to application severities defined by the enum `VCSAgAppSev` and the debug severity macros map to severities defined by the enum `VCSAgDbgSev`. For example, in the `VCSAgApiDefs.h` header file, these enumerated types are defined as:

```
enum VCSAgAppSev {
    AG_CRITICAL,
    AG_ERROR,
    AG_WARNING,
    AG_NOTICE,
    AG_INFORMATION
};

enum VCSAgDbgSev {
    DBG1,
    DBG2,
    DBG3,
    .
    .
    DBG21,
    DBG_SEV_End
};
```

With the severity macros, agent developers need not specify the name of the function, the file name, and the line number in each log call. The name of the function, however, must be initialized by using the macro `VCSAG_LOG_INIT`. See “[Initializing function_name using VCSAG_LOG_INIT](#)” on page 112.

Initializing `function_name` using `VCSAG_LOG_INIT`

One requirement for logging of messages included in C++ functions is to initialize the *function_name* variable within *each* function. The macro, `VCSAG_LOG_INIT`, defines a local constant character string to store the function name:

```
VCSAG_LOG_INIT(func_name) const char *_function_name_ =  
func_name
```

For example, the function named “res_offline” would contain:

```
void res_offline (int a, char *b)  
{  
    VCSAG_LOG_INIT("res_offline");  
    .  
    .  
}
```

Note: If the function name is not initialized with the `VCSAG_LOG_INIT` macro, when the agent is compiled, errors indicate that the name of the function is not defined.

See the “[Examples of logging APIs used in a C++ agent](#)” on page 114 for more examples of the `VCSAG_LOG_INIT` macro.

Log category

The log category for the agent is defined using the primitive `VCSAgSetLogCategory (cat_ID)` within the `VCSAgStartup` entry point. In the following example, the log category is set to 10051:

```
VCSEXPORT void VCDECL VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor,
res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline,
res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline,
res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
    VCSAgSetLogCategory(10051);

    char *s = setlocale(LC_ALL, NULL);
    VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, "Locale is
        %s", s);
}
```

You do not need to set the log category for debug messages, which is 50 by default.

Examples of logging APIs used in a C++ agent

```
#include <stdio.h>
#include <locale.h>
#include "VCSAgApi.h"

void res_attr_changed(const char *res_name, const char
    *changed_res_name, const char *changed_attr_name, void
    **new_val)
{
    /*
     * NOT REQUIRED if the function is empty or is not logging
     * any messages to the agent log file
     */
    VCSAG_LOG_INIT("res_attr_changed");
}

extern "C" unsigned int
res_clean(const char *res_name, VCSAgWhyClean wc, void
    **attr_val)
{
    VCSAG_LOG_INIT("res_clean");
    if ((attr_val) && (*attr_val)) {
        if ((remove((CHAR *)(*attr_val)) == 0) || (errno
            == ENOENT)) { return 0;          // Success
        }
    }
    return 1;          // Failure
}

void res_close(const char *res_name, void **attr_val)
{
    VCSAG_LOG_INIT("res_close");
}
//
// Determine if the given file is online (file exists) or
// offline (file does not exist).
//
extern "C" VCSAgResState
res_monitor(const char *res_name, void **attr_val, int
    *conf_level)
{
    VCSAG_LOG_INIT("res_monitor");

    VCSAgResState state = VCSAgResUnknown;
    *conf_level = 0;

    /*
     * This msg will be printed for all resources if VCS_DBG4
     * is enabled for the resource type. Else it will be
     * logged only for that resource that has the dbg level
     * VCS_DBG4 enabled
     */
}
```

```
VCSAG_RES_LOG_MSG(VCS_DBG4, VCS_DEFAULT_FLAGS, "PathName
    is(%s)", (CHAR *)(*attr_val));

if ((attr_val) && (*attr_val)) {
    struct stat stat_buf;
    if ( (stat((CHAR *)(* attr_val), &stat_buf) == 0)
        && (strlen((CHAR *)(* attr_val)) != 0) ) {
        state = VCSAgResOnline; *conf_level = 100;

    }
    else {

        state = VCSAgResOffline;
        *conf_level = 0;

    }

}
VCSAG_RES_LOG_MSG(VCS_DBG7, VCS_DEFAULT_FLAGS, "State is
    (%d)", (int)state);
return state;
}
extern "C" unsigned int
res_online(const char *res_name, void **attr_val) {
    int fd = -1;
    VCSAG_LOG_INIT("res_online");
    if ((attr_val) && (*attr_val)) {
        if (strlen((CHAR *)(* attr_val)) == 0) {
            VCSAG_LOG_MSG(VCS_WARNING, 3001, VCS_DEFAULT_FLAGS,
                "The value for PathName attribute is not
                specified");

            VCSAG_CONSOLE_LOG_MSG(VCS_WARNING, 3001,
                VCS_DEFAULT_FLAGS,
                "The value for PathName attribute is not
                specified");

            return 0;
        }
    }
    if (fd = creat((CHAR *)(*attr_val), S_IRUSR|S_IWUSR) < 0) {

        VCSAG_LOG_MSG(VCS_ERROR, 3002, VCS_DEFAULT_FLAGS,
            "Resource could not be brought up because, "
            "the attempt to create the file(%s) failed "
            "with error(%d)", (CHAR *)(*attr_val), errno);
    }
}
```

```

        VCSAG_CONSOLE_LOG_MSG(VCS_ERROR, 3002,
            VCS_DEFAULT_FLAGS,
            "Resource could not be brought up because, "
            "the attempt to create the file(%s) failed "
            "with error(%d)", (CHAR *)(*attr_val), errno);
        return 0;
    }

    close(fd);
}

return 0;
}

extern "C" unsigned int
res_offline(const char *res_name, void **attr_val)
{
    VCSAG_LOG_INIT("res_offline");
    if ((attr_val) && (*attr_val) && (remove((CHAR*)
        (*attr_val)) != 0) && (errno != ENOENT)) {

        VCSAG_LOG_MSG(VCS_ERROR, 14002, VCS_DEFAULT_FLAGS,
            "Resource could not be brought down because, the
            attempt to remove the file(%s) failed with
            error(%d)", (CHAR *)(*attr_val), errno);

        VCSAG_CONSOLE_LOG_MSG(VCS_ERROR, 14002,
            VCS_DEFAULT_FLAGS, "Resource could not be brought
            down because, the attempt to remove the file(%s)
            failed with error(%d)", (CHAR *)(*attr_val), errno);
    }

    return 0;
}

```

```

void res_open(const char *res_name, void **attr_val)
{
    VCSAG_LOG_INIT("res_open");
}

VCSEXPORT void VCSDECL VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor,
res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline,
res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline,
res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);

    VCSAgSetLogCategory(2001);

    char *s = setlocale(LC_ALL, NULL);
    VCSAG_LOGDBG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, "Locale is
    %s", s);
}

```

Script entry point logging functions

For script based entry points, use the functions described in this section for message logging purposes.

Note: Symantec recommends that you do not use the `halog` command in script entry points.

The logging functions are available in the `ag_i18n_inc` module.

VCSAG_SET_ENVS: See “[VCSAG_SET_ENVS](#)” on page 118.

VCSAG_LOG_MSG: See “[VCSAG_LOG_MSG](#)” on page 121.

VCSAG_LOGDBG_MSG: See “[VCSAG_LOGDBG_MSG](#)” on page 122.

VCSAG_CONSOLE_LOG_MSG: Logs a message to the `vcsonclientd` log file.

Using functions in scripts

The script-based entry points require a line that specifies the file defining the logging functions. Include the following line exactly once in each script. The line should precede the use of any of the log functions.

- Shell Script include file

```
. ${VCSONE_HOME:-/opt/VRTSvcs}/bin/ag_i18n_inc.sh
```
- Perl Script include file

```
use ag_i18n_inc;
```

VCSAG_SET_ENVS

The VCSAG_SET_ENVS function is used in each script-based entry point file. Its purpose is to set and export environment variables that identify the agent’s category ID, the agent’s name, the resource’s name, and the entry point’s name. With this information set up in the form of environment variables, the logging functions can handle messages and their arguments in the unified logging format without repetition within the scripts.

The VCSAG_SET_ENVS function sets the following environment variables for a resource:

VCSAG_LOG_CATEGORY	Sets the category ID. For custom agents, Symantec assigns the category ID. See the category ID description in “ UMI ” on page 107. NOTE: For bundled agents, the category ID is pre-assigned, based on the platform (Solaris, Linux, AIX, HP-UX, or Windows) for which the agent is written.
VCSAG_LOG_AGENT_NAME	The absolute path to the agent. For example: <pre>UNIX: /opt/VRTSvcs/bin/<i>resource_type</i> Windows: %VCSONE_HOME%\bin\<i>resource_type</i></pre> Since the entry points are invoked using their absolute paths, this environment variable is set at invocation. If the agent developer wishes, this agent name can also be hard coded and passed as an argument to the VCSAG_SET_ENVS function

VCSAG_LOG_SCRIPT_NAME	The absolute path to the entry point script. For example: UNIX: /opt/VRTSvcs/bin/<i>resource_type</i>/online Windows: C:\Program Files\Veritas\Cluster Server One\bin\resource_type\online Since the entry points are invoked using their absolute paths, this environment variable is set at invocation. The script name variable is can be overridden.
VCSAG_LOG_RESOURCE_NAME	The resource is specified in the call within the entry point: VCSAG_SET_ENVS <i>\$resource_name</i>

VCSAG_SET_ENVS examples, Shell script entry points

The VCSAG_SET_ENVS function must be called before any of the other logging functions.

- A minimal call:

```
VCSAG_SET_ENVS ${resource_name}
```

- Setting the category ID:

```
VCSAG_SET_ENVS ${resource_name} ${category_ID}
VCSAG_SET_ENVS ${resource_name} 1062
```

- Overriding the default script name:

```
VCSAG_SET_ENVS ${resource_name} ${script_name}
VCSAG_SET_ENVS ${resource_name} "monitor"
```

- Setting the category ID and overriding the script name:

```
VCSAG_SET_ENVS ${resource_name} ${script_name}
${category_id}
VCSAG_SET_ENVS ${resource_name} "monitor" 1062
```

Or,

```
VCSAG_SET_ENVS ${resource_name} ${category_id}
${script_name}
VCSAG_SET_ENVS ${resource_name} 1062 "monitor"
```

VCSAG_SET_ENVS examples, Perl script entry points

- A minimal call:

```
VCSAG_SET_ENVS ($resource_name);
```

- Setting the category ID:

```
VCSAG_SET_ENVS ($resource_name, $category_ID);
VCSAG_SET_ENVS ($resource_name, 1062);
```

- Overriding the script name:

```
VCSAG_SET_ENVS ($resource_name, $script_name);
VCSAG_SET_ENVS ($resource_name, "monitor");
```

- Setting the category ID and overriding the script name:

```
VCSAG_SET_ENVS ($resource_name, $script_name, $category_id);
VCSAG_SET_ENVS ($resource_name, "monitor", 1062);
```

Or,

```
VCSAG_SET_ENVS ($resource_name, $category_id, $script_name);
VCSAG_SET_ENVS ($resource_name, 1062, "monitor");
```

VCSAG_LOG_MSG

The VCSAG_LOG_MSG function can be used to pass normal agent messages to the halog utility. At a minimum, the function must include the severity, the message within quotes, and a message ID. Optionally, the function can also include parameters and specify an encoding format.

The msg is logged in the vcsoneclientd log. If the logged msg sev is CRITICAL, ERROR, the msg is logged in the engine log also.

Severity Levels (<i>sev</i>)	"C" - critical, "E" - error, "W" - warning, "N" - notice, "I" - information; place error code in quotes
Message (<i>msg</i>)	A text message within quotes; for example: "One file copied"
Message ID (<i>msgid</i>)	An integer between 0 and 65535
Encoding Format	UTF-8, ASCII, or UCS-2 in the form: "-encoding <i>format</i> "
Parameters	Parameters (up to six), each within quotes

VCSAG_LOG_MSG examples, Shell script entry points

- Calling a function without parameters or encoding format:

```
VCSAG_LOG_MSG "<sev>" "<msg>" <msgid>
VCSAG_LOG_MSG "C" "Two files found" 140
```

- Calling a function with one parameter, but without encoding format:

```
VCSAG_LOG_MSG "<sev>" "<msg>" <msgid> "<param1>"
VCSAG_LOG_MSG "C" "$count files found" 140 "$count"
```

- Calling a function with a parameter and encoding format:

```
VCSAG_LOG_MSG "<sev>" "<msg>" <msgid> "-encoding <format>"
"<param1>"
VCSAG_LOG_MSG "C" "$count files found" 140 "-encoding utf8"
"$count"
```

Note that if encoding format and parameters are passed to the functions, the encoding format must be passed before any parameters.

VCSAG_LOG_MSG examples, Perl script entry points

- Calling a function without parameters or encoding format:

```
VCSAG_LOG_MSG (<sev>, <msg>, <msgid>);
VCSAG_LOG_MSG ("C", "Two files found", 140);
```
- Calling a function with one parameter, but without encoding format:

```
VCSAG_LOG_MSG (<sev>, <msg>, <msgid>, <param1>);
VCSAG_LOG_MSG ("C", "$count files found", 140, "$count");
```
- Calling a function with one parameter and encoding format:

```
VCSAG_LOG_MSG (<sev>, <msg>, <msgid>, "-encoding
<format>", <param1>);
VCSAG_LOG_MSG ("C", "$count files found", 140, "-encoding
utf8", "$count");
```

Note that if encoding format and parameters are passed to the functions, the encoding format must be passed before any parameters.

VCSAG_LOGDBG_MSG

This function can be used to pass debug messages to the `halog` utility. At a minimum, the severity must be indicated along with a message. Optionally, the encoding format and parameters may be specified.

Severity (<i>dbg</i>)	An integer indicating a severity level, 1 to 21.
Message (<i>msg</i>)	A text message in quotes; for example: "One file copied"
Encoding <i>Format</i>	UTF-8, ASCII, or UCS-2 in the form: "-encoding <i>format</i> "
Parameters	Parameters (up to six), each within quotes

VCSAG_LOGDBG_MSG examples, Shell script entry points

- Calling a function without encoding or parameters:

```
VCSAG_LOGDBG_MSG <dbg> <msg>
VCSAG_LOGDBG_MSG 1 "This is string number 1"
```
- Calling a function with a parameter, but without encoding format:

```
VCSAG_LOGDBG_MSG <dbg> <msg> <param1>
VCSAG_LOGDBG_MSG 2 "This is string number $count" "$count"
```
- Calling a function with a parameter and encoding format:

```
VCSAG_LOGDBG_MSG <dbg> <msg> "-encoding <format>" "$count"
VCSAG_LOGDBG_MSG 2 "This is string number $count" "$count"
```

VCSAG_LOGDBG_MSG examples, Perl script entry points

- Calling a function:

```
VCSAG_LOGDBG_MSG (<dbg>, "<msg>");
VCSAG_LOGDBG_MSG (1 "This is string number 1");
```

- Calling a function with a parameter, but without encoding format:

```
VCSAG_LOGDBG_MSG (<dbg>, "<msg>", "<param1>");
VCSAG_LOGDBG_MSG (2, "This is string number $count",
"$count");
```

- Calling a function with a parameter and encoding format:

```
VCSAG_LOGDBG_MSG <dbg> "<msg>" "-encoding <format>"
"<param1>"
VCSAG_LOGDBG_MSG (2, "This is string number $count",
"-encoding
utf8", "$count");
```

Example of logging functions used in a script agent

The following example shows the use of VCSAG_SET_ENVS and VCSAG_LOG_MSG functions in a shell script for the online entry point.

```
#!/bin/ksh

ResName=$1

# Parse other input arguments
:
:
VCSONE_HOME="${VCSONE_HOME:-/opt/VRTSvcs}"

. $VCSONE_HOME/bin/ag_i18n_inc.sh

# Assume the category id assigned by Symantec for this custom
agent #is 10061
VCSAG_SET_ENVS $ResName 10061

# Online entry point processing
:
:

# Successful completion of the online entry point
VCSAG_LOG_MSG "N" "online succeeded for resource $ResName" 1
"$ResName"

exit 0
```

Building a custom agent

- [Files for use in agent development](#)
- [Creating the type definition file for a custom agent](#)
- [Building a custom agent on UNIX](#)
- [Building a custom agent on Windows](#)
- [Installing the custom agent](#)
- [Defining resources for the custom resource type](#)

Files for use in agent development

The VCS One installation program provides the following files to aid agent development:

Table 6-1 Script Agents

Description	Pathname
Ready-to-use agent that includes a built-in implementation of the VCSAgStartup entry point.	UNIX: \$VCSONE_HOME/bin/Script51Agent Script51Agent cannot be used with C++ entry points.
ag_i18n_inc.pm for perl scripts ag_i18n_inc for shell scripts.	\$VCSONE_HOME/lib /opt/VRTSvc/lib

Table 6-2 C++ Agents

Description	Pathname
Directory containing a sample C++ agent and Makefile.	UNIX: \$VCSONE_HOME/src/agent/Sample Windows: %VCSONE_HOME%\src\agent\Sample
Sample Makefile for building a C++ agent.	UNIX: \$VCSONE_HOME/src/agent/Sample/Makefile Windows: %VCSONE_HOME%\src\agent\Sample
Entry point templates for C++ agents.	UNIX: \$VCSONE_HOME/src/agent/Sample/agent.C Windows: %VCSONE_HOME%\src\agent\Sample.C

Creating the type definition file for a custom agent

The agent you create requires a resource type definition file. This file performs the function of providing a general type definition of the resource and its unique attributes.

Naming convention for the type definition file

Name the resource type definition file following the convention

resource_typeTypes.platform.xml.

- *resource_type* represents the name of the custom type. For example, for the Oracle agent, *resource_type* could assume the value Oracle.
- *platform* represents the operating system and could assume any of the following values: aix, linux, sun, win

For example, for the resource type XYZ on Solaris, the file would be *XYZTypes.sun.xml*.

For example, for the resource type XYZ on Windows, the file would be *XYZTypes.win.xml*.

Example: Type definition file for the FileOnOff agent on Linux

An example XML file for the FileOnOff resource to run on the Linux platform follows:

```
<!DOCTYPE config SYSTEM "types.dtd">
<config>

<agent name="FileOnOff" platform="linux/x86" version="1.0">
...
<!--Attributes list for this agent-->

<attributes>

    <attribute name="ArgList" type="str" dimension="vector">
        <attr_description>" "</attr_description>
        <default>
            <val>PathName</val>
        </default>
    </attribute>

    <attribute name="PathName" type="str" dimension="scalar">
        <display_name>"PathName"</display_name>
        <must_configure>1</must_configure>
        <unique>1</unique>
        <important>1</important>
```

```

        <attr_description>Specifies the complete pathname,
            starting with the slash (/) preceding the file
            name. </attr_description>
    </attribute>

<!--List of file installed by this agent-->
<agentfiles>
    <file name="VCSONE_HOME/bin/FileOnOf/FileOnOffAgent" />
</agentfiles>
</agent>
</config>

```

Example: Type definition file for the FileOnOff agent on Windows

An example XML file for the FileOnOff resource to run on the Windows platform follows:

```

<!DOCTYPE config SYSTEM "types.dtd">
<config>

<agent name="FileOnOff" platform="windows/x86">
...
<!--Attributes list for this agent-->

<attributes class="static">

    <attribute name="ArgList" type="str" dimension="vector">
        <attr_description>" "</attr_description>
        <default>
            <val>PathName</val>
        </default>
    </attribute>

<attributes>
    <attribute name="TypeDefinitionVersion" type="str"
        dimension="scalar">
        <default><scalar>5.1.0.0</scalar></default>
    </attribute>
    <attribute name="MinCompatAgentVersion" type="str"
        dimension="scalar">
        <default><scalar>5.1.0.0</scalar></default>
    </attribute>

<!--List of file installed by this agent-->
<agentfiles>
    <file name="$VCSONE_HOME/bin/monitor" />
</agentfiles>
</agent>
</config>

```

Requirements for creating the agentTypes.xml file

As you examine the previous example, note the following aspects:

- A reference to the "types.dtd" document that resides in the following directory:
UNIX: \$VCSONE_HOME/conf/confxml
Windows: %VCSONE_HOME%\conf\confxml
- The name of the agent, its platform, and version
- The `ArgList` attribute, its name, type, dimension, and its values, which consist of the other attributes of the resource
- Attributes specific to the custom resource type (in this example there is only the `PathName` attribute), their names, types, dimensions, and descriptions.
- The attributes have other properties, such as `<display_name>` and `<unique>` that further define the attributes.
See the *Veritas Cluster Server One User's Guide* for more information about these properties.

Configuring attribute properties

You can configure the following properties for the attributes associated with your custom agent:

i18nstr	Can accept localized values.
parallel_local	Local when in parallel group.
temp	Temporary attribute.
static	Static attribute
important	Used to highlight the attribute in GUI.
must_configure	Attribute must be configured.
agentencrypt	Attribute value will be encrypted.

Example: Type definition for a custom agent that supports intentional offline

```
type MyCustomApp (  
    static int IntentionalOffline = 1  
    static str ArgList[] = { PathName, Arguments }  
    str PathName  
    str Arguments
```

)

Adding the custom type definition to the configuration

This section describes how to validate your XML file (custom_typeTypes.sun.xml) and add it to your configuration.

To add the type definition when the VCS One Policy Master is running

- 1 Create the following directory:
 UNIX: /tmp/custom_type
 Windows: C:\temp\custom_type
 The variable *custom_type* represents the name of the resource type you created.
- 2 Copy the custom types XML file into the newly-created directory.
- 3 Verify the XML file.

```
haconf -verify <xml_dir>
```
- 4 Run the following command:
 On UNIX:

```
haconf -xmltocmd /tmp/custom_type/custom_typeTypes.sun.xml /tmp/custom_type
```

 On Windows:

```
haconf -xmltocmd C:\temp\custom_type\custom_typeTypes.win.xml C:\temp\custom_type
```

 This converts the custom types file to a command file called with a .cmd extension in the custom_type directory. The file contains commands needed to add the type definition to the configuration.
- 5 Run the commands included in the file from the policy master:
 On UNIX:

```
/tmp/custom_type/custom_typeTypes.sun.xml.cmd
```

 On Windows:

```
C:\temp\custom_type\custom_typeTypes.sun.xml.cmd
```
- 6 Verify the custom type definition is added to the configuration:

```
hatype -display custom_type
```

To add the type definition when the VCS One Policy Master is not running

- 1 Convert the configuration in policy master database to XML:

```
haconf -dbtoxml <xml_dir>
```

 The variable *<xml_dir>* represents the policy master configuration directory, typically:
 UNIX: \$VCSONE_HOME/conf/confxml
 Windows: %VCSONE_HOME%\conf\confxml

- 2 Once you create the type definition file for the custom type, place it in the XML directory:
UNIX: `$VCSONE_HOME/conf/confxml`
Windows: `%VCSONE_HOME%\conf\confxml`
- 3 Include the custom type definition file in the main.xml file. Add the following line to the main.xml file:
`<include>custom_typeTypes.platform.xml</include>`
- 4 Verify the updated configuration.
`haconf -verify <xml_dir>`
- 5 Load the XML configuration to the policy master database.
`haconf -loadadb <xml_dir>`

Building a custom agent on UNIX

The following sections describe different ways to build an agent, using the “FileOnOff” resource as an example. For test purposes, instructions for installing the agent on a single system are also provided.

The examples use the keyword `VCSONE_HOME` to denote the VCS One installation directory. Substitute `VCSONE_HOME` with the actual directory name when entering commands.

The examples assume:

- VCS One is installed under `/opt/VRTSvcs` by default. If your installation directory is different, change `VCSONE_HOME` accordingly.
- You have created a FileOnOff type definition file.
See [“Creating the type definition file for a custom agent”](#) on page 127.

Note the following about the FileOnOff agent entry points. A FileOnOff resource represents a regular file.

- The FileOnOff `online` entry point creates the file if it does not already exist.
- The FileOnOff `offline` entry point deletes the file.
- The FileOnOff `monitor` entry point returns online and confidence level 100 if the file exists; otherwise, it returns offline.

Example: Using script entry points on UNIX

The following example shows how to build the FileOnOff agent without writing and compiling any C++ code. This example implements the `online`, `offline`, and `monitor` entry points only.

Example: implementing entry points using scripts

- 1 Create the directory `/opt/VRTSvcs/bin/FileOnOff`:

```
mkdir /opt/VRTSvcs/bin/FileOnOff
```
- 2 Use the VCS One agent `/opt/VRTSvcs/bin/Script51Agent` as the FileOnOff agent. Copy this file to the following path:
`/opt/VRTSvcs/bin/FileOnOff/FileOnOffAgent`
or create a link.
To copy the agent binary:

```
cp /opt/VRTSvcs/bin/Script51Agent
/opt/VRTSvcs/bin/FileOnOff/FileOnOffAgent
```

To create a link to the agent binary:

```
ln -s /opt/VRTSvcs/bin/Script51Agent
/opt/VRTSvcs/bin/FileOnOff/FileOnOffAgent
```
- 3 Implement the `online`, `offline`, and `monitor` entry points using scripts. Use any editor.
 - Create the file `/opt/VRTSvcs/bin/FileOnOff/online` with the contents:

```
# !/bin/sh
# Create the file specified by the PathName
# attribute.
touch $2
exit 0
```
 - Create the file `/opt/VRTSvcs/bin/FileOnOff/offline` with the contents:

```
# !/bin/sh
# Remove the file specified by the PathName
# attribute.
rm $2
exit 0
```
 - Create the file `/opt/VRTSvcs/bin/FileOnOff/monitor` with the contents:

```
# !/bin/sh
# Verify file specified by the PathName attribute
# exists.
if test -f $2
then exit 110;
else exit 100;
fi
```

- 4 Additionally, you can implement the `info` and `action` entry points. For the `action` entry point, create a subdirectory named “actions” under the `agent` directory, and create scripts with the same names as the `action_tokens` within the subdirectory.

Example: Using VCSAgStartup() and script entry points on UNIX

The following example shows how to build the FileOnOff agent using your own VCSAgStartup entry point. This example implements the VCSAgStartup, online, offline, and monitor entry points only.

Example: implementing agent using VCSAgStartup and script entry points

- 1 Create the following directory:

```
mkdir /opt/VRTSvcs/src/agent/FileOnOff
```
- 2 Copy the contents from the sample agent directory to the directory you created in the previous step:

```
cp /opt/VRTSvcs/src/agent/Sample/*  
/opt/VRTSvcs/src/agent/FileOnOff
```
- 3 Change to the new directory:

```
cd /opt/VRTSvcs/src/agent/FileOnOff
```
- 4 Edit the file `agent.C` and modify the `VCSAgStartup()` function (the last several lines) to match the following example:

```
void VCSAgStartup() {  
    VCSAgInitEntryPointStruct(V51);  
  
    // Do not configure any entry points because  
    // this example does not implement any of them  
    // using C++.  
  
    VCSAgSetLogCategory(10041);  
}
```
- 5 Compile `agent.C` and build the agent by invoking `make`. (Makefile is provided.)

```
make
```
- 6 Create a directory for the agent:

```
mkdir /opt/VRTSvcs/bin/FileOnOff
```
- 7 Install the FileOnOff agent.

```
make install AGENT=FileOnOff
```
- 8 Implement the online, offline, and monitor entry points.
See [“Example: Using script entry points on UNIX”](#) on page 132.

Implementing entry points using C++

To implement entry points using C++

- 1 Edit `agent.C` to customize the implementation; `agent.C` is located in the directory `$VCSONE_HOME/src/agent/Sample`.
- 2 After completing the changes to `agent.C`, invoke the `make` command to build the agent. The command is invoked from `$VCSONE_HOME/src/agent/Sample`, where the `Makefile` is located.
- 3 Name the agent binary: `resource_typeAgent`.
- 4 Place the agent in the directory `$VCSONE_HOME/bin/resource_type`. For example, the agent binary for Oracle would be `$VCSONE_HOME/bin/Oracle/OracleAgent`.

Example: Using C++ entry points on UNIX

The example in this section shows how to build the `FileOnOff` agent using your own `VCSAgStartup` entry point and the C++ version of `online`, `offline`, and `monitor` entry points. This example implements the `VCSAgStartup`, `online`, `offline`, and `monitor` entry points only.

Example: VCSAgStartup and C++ entry points

- 1 Edit the file `agent.C` and modify the `VCSAgStartup()` function (the last several lines) to match the following example:

```
// Description: This functions registers the entry points //
void VCSAgStartup()
{
    VCSAG_LOG_INIT("VCSAgStartup");

    VCSAgSetLogCategory(10051);
    VCSAgInitEntryPointStruct(V51);

    VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOnline, res_online);
    VCSAgValidateAndSetEntryPoint(VCSAgEPOffline, res_offline);
    VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}
```

- 2 Modify `res_online()`:

```
// This is a C++ implementation of the online entry
// point for the FileOnOff resource type. This function
// brings online a FileOnOff resource by creating the
// corresponding file. It is assumed that the complete
// pathname of the file will be passed as the first
// ArgList attribute.
```

```

unsigned int res_online(const char *res_name, void **attr_val) {
    WCHAR      **new_args = NULL;
    DWORD      ret;
    HANDLE      hFile = NULL;
    wchar_t     *pathName = NULL;

    VCSAG_LOG_INIT("res_online");
    VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, I, "Inside
        online.");
    ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val, VCSAgUCS2,
        (void ***)&new_args);
    if (ret)
    {
        VCSAG_LOG_MSG(VCS_NOTICE, 1001, VCS_DEFAULT_FLAGS,
            I, "Unable to get the arguments");
        goto exit;
    }

    int index_of_attr = -1;
    ret = vcsag_get_attr_value((wchar_t**)new_args, I, "PathName",
        &pathName, index_of_attr, 1);
    if (ret != VCSAG_SUCCESS)
    {
        goto exit;
    }

    hFile = CreateFile(pathName,
        GENERIC_READ | GENERIC_WRITE, 0, NULL,
        OPEN_ALWAYS,
        res_ATTRIBUTE_NORMAL, (HANDLE)NULL);

    if (!hFile || hFile == INVALID_HANDLE_VALUE)
    {
        VCSAG_LOG_MSG(VCS_ERROR, 1002, VCS_DEFAULT_FLAGS,
            I, "Unable to create the file"
        );
    }
    else
    {
        VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, I, "Online
            successful.");
        CloseHandle(hFile);
    }

exit:
    if (new_args)
        VCSAgDelEncodedArgList((void**)new_args);
    return 0;
}

```

3 Modify res_offline():

```

// Function:      res_offline
// Description: This function deletes the file //

unsigned int res_offline(const char *res_name, void **attr_val)
{
    WCHAR      **new_args = NULL;
    DWORD      ret;
    wchar_t     *pathName = NULL;

    VCSAG_LOG_INIT("res_offline");
    VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,I,"Inside
        offline.");
    ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val, VCSAgUCS2,
        (void ***)&new_args);

    if (ret)
    {
        VCSAG_LOG_MSG(VCS_NOTICE, 1001, VCS_DEFAULT_FLAGS,
            I,"Unable to get the arguments");
        goto exit;
    }

    int index_of_attr = -1;
    ret = vcsag_get_attr_value((wchar_t**)new_args, I,"PathName",
        &pathName, index_of_attr, 1);
    if(ret !=VCSAG_SUCCESS)
    {
        goto exit;
    }

    if (!DeleteFile(pathName))
    {
        VCSAG_LOG_MSG(VCS_ERROR, 1003, VCS_DEFAULT_FLAGS,
            I,"Unable to delete the file");
    }
    else
    {
        VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,I,"offline
            successful.");
    }

    exit:
    if(new_args)
        VCSAgDelEncodedArgList((void**)new_args);
    return 0;
}

```

- 4 Modify the `res_monitor()`, function.
See [“Example: Using C++ and script entry points on UNIX”](#) on page 138.
- 5 Compile agent.C and build the agent by invoking make. (Makefile is provided.)

- ```
make
```
- 6 Create the directory for the agent binaries:  

```
mkdir /opt/VRTSvcs/bin/FileOnOff
```
  - 7 Install the FileOnOff agent.  

```
make install AGENT=FileOnOff
```

## Example: Using C++ and script entry points on UNIX

The following example shows how to build the FileOnOff agent using your own VCSAgStartup entry point, the C++ version of the monitor entry point, and script versions of online and offline entry points. This example implements the VCSAgStartup, online, offline, and monitor entry points only.

### Example: implementing agent using VCSAgStartup, C++, and script entry points

- 1 Create a directory for the agent:  

```
mkdir /opt/VRTSvcs/src/agent/FileOnOff
```
- 2 Copy the contents from the sample agent to the directory you created in the previous step:  

```
cp /opt/VRTSvcs/src/agent/Sample/*
 /opt/VRTSvcs/src/agent/FileOnOff
```
- 3 Change to the new directory:  

```
cd /opt/VRTSvcs/src/agent/FileOnOff
```
- 4 Edit the file agent.C and modify the VCSAgStartup() function (the last several lines) to match the following example:

```
// Description: This functions registers the entry points //
void VCSAgStartup()
{
 VCSAG_LOG_INIT("VCSAgStartup");

 VCSAgSetLogCategory(10051);
 VCSAgInitEntryPointStruct(V51);

 VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor, res_monitor);
 VCSAgValidateAndSetEntryPoint(VCSAgEPClean, res_clean);
}
```

- 5 Modify the res\_monitor() function:

```
// Function: res_monitor
// Description: Determine if the given file is online (file exists)
// or offline (file does not exist).

VCSAgResState res_monitor(const char *res_name, void
 **attr_val, int *conf_level)
```

```

{
 VCSAgResState state;
 WCHAR **new_args = NULL;
 DWORD ret, attrs;
 wchar_t *pathName = NULL;

 VCSAG_LOG_INIT("res_monitor");
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, I, "Inside
 monitor.");
 ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val, VCSAgUCS2,
 (void ***)&new_args);
 if (ret)
 {
 VCSAG_LOG_MSG(VCS_NOTICE, 1001, VCS_DEFAULT_FLAGS,
 I, "Unable to get the arguments");
 state = VCSAgResUnknown;
 goto exit;
 }

 state = VCSAgResUnknown;
 int index_of_attr = -1;
 ret = vcsag_get_attr_value((wchar_t**)new_args,
I, "PathName",
 &pathName, index_of_attr, 1);

 if (ret != VCSAG_SUCCESS)
 {
 goto exit;
 }

 attrs = GetFileAttributes(pathName);
 if (attrs == 0xffffffff)
 {
 state = VCSAgResOffline;
 *conf_level = 0;
 }
 else
 {
 state = VCSAgResOnline;
 *conf_level = 100;
 }

 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, I, "Monitor
 successful.");
exit:
 if (new_args)
 VCSAgDelEncodedArgList((void**)new_args);
 return state;
}

```

- 6** Compile `agent.C` and build the agent by invoking `make`. (Makefile is provided.)

```
make
```

- 7 Build the script entry points for the agent.  
See “[Example: Using script entry points on UNIX](#)” on page 132
- 8 Create a directory for the agent:  

```
mkdir /opt/VRTSvcs/bin/FileOnOff
```
- 9 Install the FileOnOff agent.

## Building a custom agent on Windows

The following sections describe different ways to build an agent, using the “FileOnOff” resource as an example. For test purposes, instructions for installing the agent on a single system are also provided.

The examples use the keyword `VCSONE_HOME` to denote the VCS One installation directory. Substitute `VCSONE_HOME` with the actual directory name when entering commands.

You have created a FileOnOff type definition file.

See “[Creating the type definition file for a custom agent](#)” on page 127.

Note the following about the FileOnOff agent entry points. A FileOnOff resource represents a regular file.

- The FileOnOff `online` entry point creates the file if it does not already exist.
- The FileOnOff `offline` entry point deletes the file.
- The FileOnOff `monitor` entry point returns online and confidence level 100 if the file exists; otherwise, it returns offline.

## Recommendations for building custom agents on Windows

- Symantec recommends naming the agent DLL `resource_type.dll`. Place the agent in `%VCSONE_HOME%\bin\resource_type`. The agent DLL for Oracle would be `%VCSONE_HOME%\bin\Oracle\Oracle.dll`.
- If all entry points are scripts, all scripts should be in the directory `%VCSONE_HOME%\bin\resource_type`. Copy the Script51Agent into the agent directory as `%VCSONE_HOME%\bin\resource_type\resource_typeAgent`.
- If entry points are implemented using scripts, the script file must be placed in the directory `%VCSONE_HOME%\bin\resource_type`. It must be named in the following manner:  
`%VCSONE_HOME%\bin\resource_type\entry_point.extn`  
The extension (`.extn`) assumes the value `exe` for executable programs, `sh` for shell scripts compatible with MKS Toolkit, `bat` or `cmd` for batch files, `pl` for Perl scripts, `js` for Javascript, or `vbs` for Visual Basic scripts.  
For example, if the online entry point for Oracle is implemented using Perl, the online script must be: `%VCSONE_HOME%\bin\Oracle\online.pl`.
- For multi-system configurations, you must install the agent on each system.
- Compiling is not required if all entry points are implemented using scripts. On Windows, a copy of `Default50Agent.dll` is sufficient.

Compiling is required to build the agent if any entry points are implemented using C++.

## Example: Using C++ entry points on Windows

The example in this section shows how to build the FileOnOff agent using your own VCSAgStartup entry point and the C++ version of online, offline, and monitor entry points. This example implements the VCSAgStartup, online, offline, and monitor entry points only.

---

**Note:** You must first install Visual C++ on the system on which the agent will be built.

---

### Example: VCSAgStartup and C++ entry points

- 1 Create the following directory for the agent sources:  
%VCSONE\_HOME%\src\agent\FileOnOff:  

```
C:\> md %VCSONE_HOME%\src\agent\FileOnOff
```
- 2 Copy the contents of the sample agent directory to the directory you created in the previous step:  

```
C:\> cp %VCSONE_HOME%\src\agent\Sample* %VCSONE_HOME%\src\agent\FileOnOff
```
- 3 Change to the new directory:  

```
C:\> cd %VCSONE_HOME%\src\agent\FileOnOff
```
- 4 Edit the file sample.C and modify the VCSAgStartup() function (the last several lines) to match the following example:

```
void VCSAgStartup()
{
 // This example implements online, offline, and
 // monitor entry points using C++.

 VCSAG_LOG_INIT("VCSAgStartup");
 VCSAgSetLogCategory(10051);

 VCSAgInitEntryPointStruct(V51);
 VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor,
res_monitor);
 VCSAgValidateAndSetEntryPoint(VCSAgEPOnline,
res_online);
 VCSAgValidateAndSetEntryPoint(VCSAgEPOffline,
res_offline);
}
```

5 Modify `res_online()` and `res_offline()`:

```

// This is a C++ implementation of the online entry
// point for the FileOnOff resource type.
// This function brings online a FileOnOff
// resource by creating the corresponding file.
// It is assumed that the complete pathname of the
// file will be passed as the first
// ArgList attribute.

unsigned int res_online(const char *res_name,
 void **attr_val) {
 WCHAR **new_args = NULL;
 DWORD ret;
 HANDLE hFile = NULL;

 VCSAG_LOG_INIT("res_online");
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
 I, "Inside online.");

 ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val, VCSAgUCS2,
 (void ***)&new_args);

 if (ret)
 {
 VCSAG_LOG_MSG(VCS_NOTICE, 1001, VCS_DEFAULT_FLAGS,
 I, "Unable to get the arguments");
 goto exit;
 }
 hFile = CreateFile((WCHAR *)new_args[0], GENERIC_READ |
 GENERIC_WRITE, 0, NULL, OPEN_ALWAYS,
 res_ATTRIBUTE_NORMAL, (HANDLE)NULL);
 if (!hFile || hFile == INVALID_HANDLE_VALUE)
 {
 VCSAG_LOG_MSG(VCS_ERROR, 1002, VCS_DEFAULT_FLAGS,
 I, "Unable to create the file");
 }
 else
 {
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
 I, "Online successful."); CloseHandle(hFile);
 }

exit:
 if(new_args)
 VCSAgDelEncodedArgList((void**)new_args);
 return 0;
}

// This is a C++ implementation of the offline
// entry point for the FileOnOff resource type.
// This function takes offline aFileOnOff
// resource by deleting the corresponding file. It
// is assumed that the complete pathname of the

```

```

 // file will be passed as the first
 // ArgList attribute.
unsigned int res_offline(const char *res_name,
void **attr_val)
{
 WCHAR **new_args = NULL;
 DWORD ret;

 VCSAG_LOG_INIT("res_offline");
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
 I,"Inside offline.");

 ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val,
 VCSAgUCS2, (void ***)&new_args);

 if (ret)
 {
 VCSAG_LOG_MSG(VCS_NOTICE, 1001, VCS_DEFAULT_FLAGS,
 I,"Unable to get the arguments");
 goto exit;
 }

 if (!DeleteFile((WCHAR *)new_args[0]))
 {
 VCSAG_LOG_MSG(VCS_ERROR, 1003, VCS_DEFAULT_FLAGS,
 I,"Unable to delete the file");
 }
 else
 {
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
 I,"offline successful.");
 }

exit:
 if(new_args)
 VCSAgDelEncodedArgList((void**)new_args);
 return 0;
}

```

- 6 Modify the `res_monitor()` function.  
See “[Example: Using VCSAgStartup\(\) and script entry points on Windows](#)” on page 146.
- 7 Compile `sample.c` and build the agent by invoking `nmake`. (Makefile is provided.)  
C:\> **nmake**
- 8 Create a directory for the agent binaries:  
%VCSONE\_HOME%\bin\FileOnOff:  
C:\> **md %VCSONE\_HOME%\bin\FileOnOff**
- 9 Place the `sample.dll` as the FileOnOff agent library.

```
C:\> copy sample.dll %VCSONE_HOME%\bin\FileOnOff\FileOnOff.dll
```

## Implementing entry points using scripts

If entry points are implemented using scripts, the script file must be placed in the directory %VCSONE\_HOME%\bin\resource\_type. It must be named in the following manner:

```
%VCSONE_HOME%\bin\resource_type\entry_point.extn
```

The extension (.extn) assumes the value exe for executable programs, sh for shell scripts compatible with MKS Toolkit, bat or cmd for batch files, pl for Perl scripts, js for Javascript, or vbs for Visual Basic scripts.

For example, if the online entry point for Oracle is implemented using Perl, the online script must be: %VCSONE\_HOME%\bin\Oracle\online.pl.

We also recommend naming the agent DLL resource\_type.dll. Place the agent in %VCSONE\_HOME%\bin\resource\_type.

## Example: Using script entry points on Windows

The following example shows how to build the FileOnOff agent without writing and compiling any C++ code. This example implements the `online`, `offline`, and `monitor` entry points only.

### Example: implementing entry points without C++

- 1 Create the directory %VCSONE\_HOME%\bin\FileOnOff:

```
C:\> md %VCSONE_HOME%\bin\FileOnOff
```

- 2 Use %VCSONE\_HOME%\bin\Default50Agent.dll as the agent DLL.

Copy the file to %VCSONE\_HOME%\bin\FileOnOff.dll, as shown below:

```
C:\> copy %VCSONE_HOME%\bin\Default50Agent.dll
 %VCSONE_HOME%\bin\FileOnOff\FileOnOff.dll
```

- 3 Implement the `online`, `offline`, and `monitor` entry points using scripts. The examples below use shell scripts. MKS Toolkit or equivalent software is required for running shell scripts. Entry points can also be implemented as batch files.

- Using any editor, create the file %VCSONE\_HOME%\FileOnOff\online.sh with the following contents:

```
Create the file specified by the PathName
attribute.
touch $2
```

- Create the file %VCSONE\_HOME%\bin\FileOnOff\offline.sh with the following contents:

```
Remove the file specified by the PathName
```

```
attribute.
rm $2
```

- Create the file %VCSONE\_HOME%\bin\FileOnOff\monitor.sh with the following contents:

```
Verify that file specified by PathName
attribute exists.
if test -f $2
then exit 110;
else exit 100;
fi
```

Additionally, you can implement the `info` and `action` entry points. For the action entry point, create a subdirectory named “actions” under the agent directory, and create scripts with the same names as the `action_tokens` within the subdirectory.

## Example: Using VCSAgStartup() and script entry points on Windows

The following example shows how to build the FileOnOff agent using your own VCSAgStartup entry point. This example implements the VCSAgStartup, online, offline, and monitor entry points only.

---

**Note:** You must first install Visual C++ on the system on which the agent will be built.

---

### Example: implementing agent using VCSAgStartup and script entry points

- 1 Create a directory for the agent source files  
%VCSONE\_HOME%\src\agent\FileOnOff:  
C:\> md %VCSONE\_HOME%\src\agent\FileOnOff
- 2 Copy the contents of the sample agent directory to the directory you created in the previous step:  
C:\> copy %VCSONE\_HOME%\src\agent\Sample\\*  
%VCSONE\_HOME%\src\agent\FileOnOff
- 3 Change to the new directory:  
C:\> cd %VCSONE\_HOME%\src\agent\FileOnOff
- 4 Edit the file sample.C and modify the VCSAgStartup() function (the last several lines) to match the following example:

```
void VCSAgStartup() {
 VCSAgInitEntryPointStruct(V51);

 // Do not configure any entry points because
 // this example does not implement any of them
 // using C++.
```

```
VCSAgSetLogCategory(10041);
}
```

- 5 Compile `sample.C` and build the agent by invoking `nmake`. (Makefile is provided.)

```
C:\> nmake
```

- 6 Create the following directory  
`%VCSONE_HOME%\bin\FileOnOff`:

```
C:\> md %VCSONE_HOME%\bin\FileOnOff
```

- 7 Place the `sample.dll` as the `FileOnOff` agent library.

```
C:\> copy sample.dll
 %VCSONE_HOME%\bin\FileOnOff\FileOnOff.dll
```

- 8 Implement the online, offline, and monitor entry points.  
 See “[Example: Using C++ entry points on Windows](#)” on page 142.

## Example: Using C++ and script entry points on Windows

The following example shows how to build the `FileOnOff` agent using your own `VCSAgStartup` entry point, the C++ version of the `monitor` entry point, and the script version of online and offline entry points. This example implements the `VCSAgStartup`, online, offline, and `monitor` entry points only.

---

**Note:** To build an agent, you must first install Visual C++ on the system on which the agent will be built.

---

### Example: implementing agent using `VCSAgStartup`, C++ and script entry points

- 1 Create a directory for the agent:

```
C:\> md %VCSONE_HOME%\src\agent\FileOnOff
```

- 2 Copy the contents of the sample agent directory to the directory you created in the previous step:

```
C:\> copy %VCSONE_HOME%\src\agent\Sample*
%VCSONE_HOME%\src\
```

- 3 Change to the new directory:

```
C:\> cd %VCSONE_HOME%\src\agent\FileOnOff
```

- 4 Edit the file `sample.C` and modify the `VCSAgStartup()` function (the last several lines) to match the following example:

```
void VCSAgStartup()
{
```

```

VCSAG_LOG_INIT("VCSAgStartup");
VCSAgSetLogCategory(10051);
VCSAgInitEntryPointStruct(V51);

VCSAgValidateAndSetEntryPoint(VCSAgEPMonitor,
res_monitor);
}

```

## 5 Modify the res\_monitor() function:

```

// This is a C++ implementation of the monitor entry
// point for the FileOnOff resource type. This function
// determines the status of a FileOnOff resource
// by checking if the corresponding file exists. It
// is assumed that the complete pathname of the file will
// be passed as the first ArgList attribute.

VCSAgResState res_monitor(const char *res_name, void
**attr_val, int *conf_level) {
 // Initialize the OUT parameters.
 VCSAgResState state;
 WCHAR **new_args = NULL;
 DWORD ret, attrs;
 VCSAG_LOG_INIT("res_monitor");
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS,
 I, "Inside monitor.");
 ret = VCSAgGetEncodedArgList(VCSAgUTF8, attr_val,
 VCSAgUCS2,
 (void ***)&new_args);
 if (ret)
 {
 VCSAG_LOG_MSG(VCS_NOTICE, 1001, VCS_DEFAULT_FLAGS,
 I, "Unable to get the arguments");
 state = VCSAgResUnknown;
 goto exit;
 }

 attrs = GetFileAttributes(new_args[0]);
 if (attrs == 0xffffffff)
 {
 state = VCSAgResOffline;
 *conf_level = 0;
 }
 else
 {
 state = VCSAgResOnline;
 *conf_level = 100;
 }
 VCSAG_RES_LOG_MSG(VCS_DBG1, VCS_DEFAULT_FLAGS, I, "Monitor
 successful.");

exit:
 if(new_args)

```

```
 VCSAgDelEncodedArgList((void**)new_args);
 return state;
}
```

- 6 Compile `sample.C` and build the agent by invoking `nmake`. (Makefile is provided.)

```
C:\> nmake
```

- 7 Create the following directory:

```
C:\Program Files\Veritas\Cluster Server One\bin\FileOnOff.
```

```
C:\> md %VCSONE_HOME%\bin\FileOnOff
```

- 8 Place the `sample.dll` as the FileOnOff agent library.

```
C:\> copy sample.dll
 %VCSONE_HOME%\bin\FileOnOff\FileOnOff.dll
```

- 9 Implement the online and offline entry points.  
See [“Example: Using C++ entry points on Windows”](#) on page 142.

## Installing the custom agent

You can install the custom agent in one of the following directories.

On UNIX:

- A user-defined directory. For example `/myagents/custom_type/`. Note that you must configure the `AgentDirectory` attribute for this option.
- `/opt/VRTSvcs/bin/custom_type/`
- `/opt/VRTSagents/ha/bin/custom_type/`

On Windows:

- A user-defined directory. For example `\myagents\custom_type\`. Note that you must configure the `AgentDirectory` attribute for this option.
- `%VCSONE_HOME%\bin\custom_type\`
- `%VCSONE_HOME%\ha\bin\custom_type/`

---

**Note:** Create the `custom_type` directory at only one of these locations.

---

Add the agent binary and the script entry points to the `custom_type` directory.

---

**Note:** To package the agent, see the documentation for the operating system. When setting up the Solaris `pkginfo` file for the installation of agents that are to run in zones, set the following variable: `SUNW_PKG_ALLZONES=true`.

---

## Defining resources for the custom resource type

When you have created a type definition for the resource and created an agent for it, you can begin to use the agent to control specific resources by adding the resources of the custom type and assigning values to resource attributes.

You can add resources and configure attributes values for the new resources using the GUI or the command-line interface.

See the *Veritas Cluster Server One User's Guide* for more information.

## Sample resource definition

In the VCS One configuration file, `main.xml`, a specific resource of the `FileOnOff` resource type may resemble:

**Snippet from a sample `main.xml`**

```

<include>types.linux.xml</include>
<include>types.sun.xml</include>
<include>types.win.xml</include> \

<group name="sg1">
 <attributes>
 ...
 </attributes>
 <resources>
 <resource name="temp_file1" type="FileOnOff">
 <attribute name="PathName"><scalar>
 "/tmp/test"</scalar>
 </attribute>
 </resource>
 </resources>
</group>

```

The include statements at the beginning of the main.xml file name the types files for each platform. The files include the FileOnOff resource type definition. The resource defined in the main.xml file specifies:

- The resource type: FileOnOff
- The name of the resource, temp\_file1
- The name of the attribute, PathName
- The value for the PathName attribute:
  - On UNIX: "/tmp/test"
  - On Windows: "C:\tmp\test"

When the resource temp\_file1 is brought online on a system by VCS One, the FileOnOff agent creates a file “test” in the specified directory on that system.

## How the FileOnOff agent uses configuration information

The information in the VCS One configuration is passed by the engine to the FileOnOff agent when the agent starts up on a node in the cluster. The information passed to the agent includes: the names of the resources of the type FileOnOff configured on the system, the corresponding resource attributes, and the values of the attributes for all of the resources of that type.

Thereafter, to bring the resource online, for example, VCS One can provide the agent with the name of the entry point (online) and the name of the resource (temp\_file01). The agent then calls the entry point and provides the resource name and values for the attributes in the ArgList to the entry point. The entry point performs its tasks.



# Testing agents

- [About testing agents](#)
- [Using debug messages](#)
- [Using the engine process to test agents](#)
- [Using the AgentServer utility to test agents](#)

## About testing agents

Before testing an agent, make sure you have built the agent and have installed and configured the agent.

## Using debug messages

You can activate agent framework debug messages by setting the value of the LogDbg attribute. This directs the framework to print messages logged with the specified severity.

See “[LogDbg](#)” on page 170.

## Debugging agent functions (entry points).

The LogDbg attribute indicates the debug severities enabled for the resource type or agent framework. Debug severities used by agent functions are in the range of DBG\_1-DBG\_21.

To enable debug logging, use the LogDbg attribute at the type-level.

To set debug severities for a particular resource-type:

```
hatype -modify <resource-type> LogDbg -add <Debug-Severity>
[<Debug-Severity> ...]
```

To remove debug severities for a particular resource-type:

```
hatype -modify <resource-type> LogDbg -delete <Debug-Severity>
[<Debug-Severity> ...]
```

To remove all debug severities:

```
hatype -modify <resource-type> LogDbg -delete -keys
```

Note that you cannot set debug severities for an individual resource.

To debug a specific agent's entry point, see the documentation for that agent. So for bundled agents, see the *Bundled Agent's Reference Guide*.

UNIX: \$VCSONE\_LOG/log/resource\_type\_A.log

Windows: %VCSONE\_HOME%\log\resource\_type\_A.txt

For example, if you want to log debug messages for the FileOnOff resource type with severity levels DBG\_3 and DBG\_4, use the `hatype` commands:

```
hatype -modify FileOnOff LogDbg -add DBG_3 DBG_4
hatype -display FileOnOff -attribute LogDbg
TYPE ATTRIBUTE VALUE
FileOnOff LogDbg DBG_3 DBG_4
```

The debug messages from the FileOnOff agent with debug severities DBG\_3 and DBG\_4 get printed to the log files. Debug messages from C++ entry points get printed to the agent log file and from script entry points will get printed to the `vcsonclientd` log file. An example line from the agent log file:

```
.
2003/06/06 11:02:35 VCS DBG_3 V-16-50-0
FileOnOff:fl.monitor:This is a debug message
FileOnOff.C:res_monitor[28]
```

## Debugging the agent framework

The LogDbg attribute indicates the debug severities enabled for the resource type or agent framework. The debug messages from the agent framework are logged with the following severities:

- **DBG\_AGDEBUG:** Enables most debug logs, which include: debugging commands received from the engine, service thread execution code path, that is when a service thread picks up a resource for running an entry point or for modification of an attribute, printing of environment variables that the agent uses, timer-related processing like sending IAmAlive messages to engine, and so on.
- **DBG\_AGINFO:** Enables debugging messages related to specific entry-point execution, including entry point exit codes, transitioning of resources between various internal-states, printing of ArgListValues before entry point invocation, values of entry-point execution related attributes like RunInContainer and PassCInfo, and so on.
- **DBG\_AGTRACE:** Enables verbose debug logging, the bulk of it being Begin and End messages for almost every function that gets called within the agent-framework, like function-tracing.

## Using the engine process to test agents

When the VCS One vcsonclientd process becomes active on a system, it automatically starts the appropriate agent processes based on the contents of the PM database.

A single agent process monitors all resources of the same type on a system.

After the VCS One vcsonclientd process is active, type the following command at the system prompt to verify that the agent has been started and is running:

```
haagent -display <resource_type>
```

For example, to test the Oracle agent, type:

```
haagent -display Oracle
```

If the Oracle agent is running, the output resembles:

| #Agent | Attribute | Value |
|--------|-----------|-------|
| Oracle | AgentFile |       |
| Oracle | Faults    | 0     |
| Oracle | Running   | Yes   |

|        |         |     |
|--------|---------|-----|
| Oracle | Started | Yes |
|--------|---------|-----|

## Test commands

The following examples show how to use commands to test the agent:

- To activate agent debug messages for C++ agents, type:  
`hatype -modify <resource_type> LogDbg -add DBG_AGINFO`
- To check the status of a resource, type:  
`hares -display <resource_name>`
- To bring a resource online, type:  
`hares -online <resource_name> -sys system`  
This causes the `online` entry point of the corresponding agent to be called.
- To take a resource offline, type:  
`hares -offline <resource_name> -sys system`  
This causes the `offline` entry point of the corresponding agent to be called.
- To deactivate agent debug messages for C++ agents, type:  
`hatype -modify <resource_type> LogDbg -delete DBG_AGINFO`

# Using the AgentServer utility to test agents

The `AgentServer` utility enables you to test agents without running the `vcsonclientd` process. The utility is part of the product package and is installed in the directory

UNIX: `$VCSONE_HOME/bin`.

Windows: `%VCSONE_HOME%\bin`

Run the `AgentServer` utility on any node when the engine process is not running.

## To start the AgentServer and access help (UNIX)

- 1 Type the following command to start `AgentServer`:

```
$VCSONE_HOME/bin/AgentServer
```

The `AgentServer` utility monitors a TCP port for messages from the agents. Configure this port by setting the following string in the `/etc/services` file to the port number.

```
vcsonetest <port number>/tcp
```

If you do not specify this value, `AgentServer` uses the default value 14142.

- 2 When `AgentServer` is started, a message prompts you to enter a command/  
For a complete list of the `AgentServer` commands, type:

```
> help
```

Output resembles:

```
The following commands are supported. (Use help for more
information on using any command.)
```

```
addattr
addres
addstaticattr
addtype
debughash
debugmemory
debugtime
delete
deleteres
modifyres
modifytype
offlineres
onlineres
print
proberes
startagent
stopagent
quit
```

- 3 For help on a specific command, type `help command_name` at the AgentServer prompt (`>`). For example, for information on how to bring a resource online, type:

```
> help onlineres
```

The output resembles:

```
Sends a message to an agent to online a resource.
Usage: onlineres <agentid> <resname>
where <agentid> is id for the agent - usually same as
the resource type name.
where <resname> is the name of the resource.
```

### To test the FileOnOff agent (UNIX)

- 1 Start the agent for the resource type:  

```
>startagent FileOnOff /opt/VRTSvcs/bin/FileOnOff/FileOnOffAgent
```

 You receive the following messages:  

```
Agent (FileOnOff) has connected.
Agent (FileOnOff) is ready to accept commands.
```
- 2 Refer to the `types.xml` and `main.xml` configuration files when testing the FileOnOff agent. In [step 3](#), the sample configuration is set up using AgentServer commands.
- 3 Complete the following steps to pass this sample configuration to the agent.

- Add a type:

```
>addtype FileOnOff FileOnOff
```

- Add attributes of the type:

```
>addattr FileOnOff FileOnOff PathName str ""
>addattr FileOnOff FileOnOff Enabled int 0
```

- Add the static attributes to the FileOnOff resource type:

```
>addstaticattr FileOnOff FileOnOff ArgList vector PathName
```

- Add the LogLevel attribute to see the debug messages from the agent:

```
>addstaticattr FileOnOff FileOnOff LogLevel str info
```

- Add a resource:

```
>addres FileOnOff file1 FileOnOff
```

- Set the resource attributes:

```
>modifyres FileOnOff file1 PathName str /tmp/VRTSvcsfile001
>modifyres FileOnOff file1 Enabled int 1
```

- 4 After adding and modifying resources, type the following command to obtain the status of a resource:

```
>proberes FileOnOff file1
```

This calls the `monitor` entry point of the FileOnOff agent.

You will receive the following messages indicating the resource status:

```
Resource(file1) is OFFLINE
Resource(file1) confidence level is 0
```

- To bring a resource online:
 

```
>onlineres FileOnOff file1
```

This calls the online entry point of the FileOnOff agent. The following message is displayed when file1 is brought online:

```
Resource(file1) is ONLINE
Resource(file1) confidence level is 100
```
  - To take a resource offline:
 

```
>offlineres FileOnOff file1
```

This calls the offline entry point of the FileOnOff agent.
- 5 View the list of agents started by the AgentServer process:
 

```
>print
```

Output resembles:

```
Following Agents are started:
FileOnOff
```
  - 6 Stop the agent:
 

```
>stopagent FileOnOff
```
  - 7 Exit from the AgentServer:
 

```
>quit
```

### To start AgentServer and test the FileOnOff agent (Windows)

- 1 Type the following command to start AgentServer:
 

```
%VCSONE_HOME%\bin\AgentServer
```

AgentServer must monitor a TCP port for messages from the agents. Configure this port by setting the following string in the file C:\WINNT\system32\drivers\etc\services:

```
vcsonetest
```

If you do not specify a value, AgentServer uses the default value 14142.
- 2 Start the agent for the resource type:
 

```
>startagent FileOnOff %VCSONE_HOME%\bin\VCSAgDriver.exe
```

. You will receive the following messages:

```
Agent (FileOnOff) has connected.
Agent (FileOnOff) is ready to accept commands.
```
- 3 Refer to the types.xml and main.xml configuration files when testing the FileOnOff agent. In the next step, the sample configuration is set up using AgentServer commands.
- 4 Complete [step n](#) through [step n](#) to pass this sample configuration to the agent.
  - Add a type:
 

```
>addtype FileOnOff FileOnOff
```
  - Add attributes of the type:

```
>addattr FileOnOff FileOnOff PathName str ""
>addattr FileOnOff FileOnOff Enabled int 0
```

- Add the static attributes to the FileOnOff resource type:

```
>addstaticattr FileOnOff FileOnOff ArgList vector PathName
```

- Add the LogLevel attribute to see the debug messages from the agent:

```
>addstaticattr FileOnOff FileOnOff LogLevel str info
```

- Add a resource:

```
>addres FileOnOff file1 FileOnOff
```

- Set the resource attributes:

```
>modifyres FileOnOff file1 PathName str
C:\\VRTSvcsfile001
>modifyres FileOnOff file1 Enabled int 1
```

- 5 After adding and modifying resources, type the following command to obtain the status of a resource:

```
>proberes FileOnOff file1
```

This calls the `monitor` entry point of the FileOnOff agent.

You will receive the following messages indicating the resource status:

```
Resource(file1) is OFFLINE
Resource(file1) confidence level is 0
```

- To bring a resource online:

```
>onlineres FileOnOff file1
```

This calls the `online` entry point of the FileOnOff agent. The following message is displayed when file1 is brought online:

```
Resource(file1) is ONLINE
Resource(file1) confidence level is 100
```

- To take a resource offline:

```
>offlineres FileOnOff file1
```

This calls the `offline` entry point of the FileOnOff agent. A status message similar to the example in [step n](#) is displayed when file1 is taken offline.

- 6 View the list of agents started by the AgentServer process:

```
>print
```

Output resembles:

```
Following Agents are started:
FileOnOff
```

- 7 Stop the agent:

```
>stopagent FileOnOff
```

- 8 Exit from the AgentServer:

```
>quit
```

# Static type attributes

- [About static attributes](#)
- [Static timeout attributes on Windows](#)
- [Static type attribute definitions](#)

## About static attributes

Predefined static resource type attributes apply to all resource types.

See “[Static type attribute definitions](#)” on page 164

When developers create agents and define the resource type definitions for them, the static type attributes become part of the type definition.

## Overriding static type attributes

Typically, the value of a static attribute of a resource type applies to all resources of the type. You can override the value of a static attribute for a specific resource without affecting the value of that attribute for other resources of that type. In this chapter, the description of each agent attribute indicates whether the attribute’s values can be overridden.

Users can override the values of static attributes two ways:

- By explicitly defining the attribute in a resource definition
- By using the `hares` command from the command line with the `-override` option

The values of the overridden attributes may be displayed using the `hares -display` command. You can remove the overridden values of static attributes by using the `hares -undo_override` option from the command line.

See the *User’s Guide* for additional information about overriding the values of static attributes.

## Static timeout attributes on Windows

Windows does not support a safe method for one thread to terminate another thread running arbitrary code. Therefore, agents on Windows wait until the entry points complete. The entry point timeouts listed in this chapter *do not* apply to Windows.

The following entry point timeouts *do not* apply for agents on Windows:

- ActionTimeout
- AttrChangedTimeout
- CloseTimeout
- CleanTimeout
- FaultOnMonitorTimeout
- InfoTimeout
- MonitorTimeout,
- OfflineTimeout
- OnlineTimeout
- OpenTimeout.

On the Windows platform, agents do not execute entry point timeouts, but wait until the entry points complete. For information concerning the timeouts listed above:

See the *User's Guide* for more information.

# Static type attribute definitions

The following sections describe the static attributes for agents.

## ActionTimeout

After the `hares -action` command has instructed the agent to perform a specified action, the `action` entry point has the time specified by the `ActionTimeout` attribute (scalar-integer) to perform the action. The value of `ActionTimeout` may be set for individual resources, if overridden.

Whether overridden or not, no matter what value is specified for `ActionTimeout`, the value is internally limited to  $0.5 * \text{MonitorInterval}$ . `MonitorInterval` attribute description is given below.

---

**Note:** This attribute is not used on Windows.

---

The default is 30 seconds. The `ActionTimeout` attribute value can be overridden.

## AEPTimeout

The `AEPTimeout` (Append Entry Point Timeout) attribute is a boolean attribute. Set this attribute to true to append the timeout value for a particular entry point to the list of arguments passed to the entry point.

For example, if you set `AEPTimeout` to 1, the agent framework passes the value of the `MonitorTimeout` attribute to the monitor entry point in the name-value tuple format. Likewise, the value of `CleanTimeout` gets passed to the clean entry point.

This feature does not apply to pre-V50 agents.

See “[About the entry point timeouts](#)” on page 40.

## AgentClass

Indicates the scheduling class for agent process.

Default is “TS”.

## AgentFailedOn

A keylist attribute indicating the systems on which the agent has failed. This is not a user defined attribute.

Default is an empty keylist.

## AgentPriority

Indicates the priority in which the agent process runs.  
Default is 0.

## AgentReplyTimeout

The `vcstoneclntd` process restarts an agent if it has not received *any* messages from the agent for the number of seconds specified by `AgentReplyTimeout`.

The default value of 130 seconds works well for most configurations. Increase this value if the `vcstoneclntd` process is restarting the agent too often during steady state of the cluster. This may occur when the system is heavily loaded or if the number of resources exceeds four hundred. Refer to the description of the command `haagent -display`. Note that the `vcstoneclntd` process will also restart a crashed agent.

The `AgentReplyTimeout` attribute value *cannot* be overridden.

---

**Note:** This attribute is not used on Windows.

---

## AgentStartTimeout

The value of `AgentStartTimeout` specifies how long the `vcstoneclntd` process waits for the initial agent “handshake” after starting the agent, before attempting to restart it.

Default is 60 seconds. The `AgentStartTimeout` attribute value *cannot* be overridden.

---

**Note:** This attribute is not used on Windows.

---

## ArgList

An ordered list of attributes whose values are passed to the open, close, online, offline, monitor, info, action, and clean entry points.

The default is an empty list. The `ArgList` attribute value cannot be overridden.

### ArgList reference attributes

Reference attributes refer to attributes of a different resource. If the value of a resource’s attribute is the name of another resource, the `ArgList` of the first resource can refer to an attribute of the second resource using the `:` operator.

For example, say, there is a type T1 whose ArgList is of the form:

```
{ Attr1, Attr2, Attr3:Attr_A }
```

where Attr1, Attr2 and Attr3 are attributes of type T1, and say for a resource res1T1 of type T1, Attr3 's value is the name of another resource, res1T2. Then the entry points for res1T1 are passed the values of attributes Attr1 and Attr2 of res1T1 and the value of attribute Attr\_A of resource res1T2.

Note that one has to first add the attribute Attr3 to type T1 before adding Attr3:Attr\_A to T1's ArgList. Only then should one modify Attr3 for a resource (res1T1) to reference another resource (res1T2). Also, the value of Attr3 can either be another resource of the same time (res2T1) or a resource of a different type (res1T2).

## AttrChangedTimeout

Maximum time (in seconds) within which the `attr_changed` entry point must complete or else be terminated. Default is 60 seconds. The `AttrChangedTimeout` attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## CleanTimeout

Maximum time (in seconds) within which the `clean` entry point must complete or else be terminated.

Default is 60 seconds. The `CleanTimeout` attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## CleanRetryLimit

Specifies the number of times the `clean` entry point needs is executed if the execution fails. The execution of `clean` entry point is considered as failed if it times out or if it exits/returns with a failure code. Once the `clean` EP is called `CleanRetryLimit` number of times, the agent framework gives up running the `Clean` entry point and sets the `ADMIN_WAIT` flag for the resource to request administrative intervention.

## CloseTimeout

Maximum time (in seconds) within which the `close` entry point must complete or else be terminated.

Default is 60 seconds. The `CloseTimeout` attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## ContainerOpts

Specifies whether the service group's load dimensions need to be passed to the resource to initialize and set CPU shares for the resource. It also specifies whether the resource needs to run in the context of a container. For each application or resource type that you want to include as part of the zone or project, you need to assign the following values to the `ContainerOpts` attribute:

- **RunInContainer (RIC)**  
RunInContainer defines whether the agent framework should run all the script-based entry points for the agent inside the container. If the attribute is set to 1, all script-based entry points are forked off inside the local container that is configured for the corresponding resource. If the attribute is set to 0, even if a resource's service group has `ContainerInfo` set, the entry point scripts for that resource will still be run in the global container.
- **PassCInfo (PCI)**  
PassCInfo specifies if you want to pass the container information, defined in the service group's `ContainerInfo` attribute, to the resource. Specify a value of 1, if you want to pass the container information to the resource. Specify a value of 0, if you do not want to pass the container information to the resource.
- **PassLoadInfo (PLI)**  
PassLoadInfo specifies if you want to pass the service groups load dimensions, defined in the service group's `Load` attribute, to the resource. The service groups load dimensions are set in the resource's `Workload` attribute. Specify a value of 1, if you want to pass the service groups load dimensions to the resource. Specify a value of 0, if you do not want to pass the service groups load dimensions to the resource.

Each resource type also contains the `Workload` attribute, which is a non-static attribute. The value for this attribute is populated from the service group's `Load` attribute, when the resource's `PassLoadInfo` (`ContainerOpts` attribute) key set to 1.

## ConfInterval

Specifies an interval in seconds. When a resource has remained online for the designated interval (all `monitor` invocations during the interval reported `ONLINE`), any earlier faults or restart attempts of that resource are ignored. This attribute is used with `ToleranceLimit` to allow the `monitor` entry point to report `OFFLINE` several times before the resource is declared `FAULTED`. If `monitor` reports `OFFLINE` more often than the number set in `ToleranceLimit`, the resource is declared `FAULTED`. However, if the resource remains online for the interval designated in `ConfInterval`, any earlier reports of `OFFLINE` are not counted against `ToleranceLimit`.

The agent framework uses the values of `MonitorInterval` (MI), `MonitorTimeout` (MT), and `ToleranceLimit` (TL) to determine how low to set the value of `ConfInterval`. The agent framework ensures that `ConfInterval` (CI) cannot be less than that expressed by the following relationship:

$$(MI + MT) * TL + MI + 10$$

Lesser specified values of `ConfInterval` are ignored. For example, assume that the values are 60 for MI, 60 for MT, and 0 for TL. If you specify any value lower than 70 for CI, the agent framework ignores the specified value and sets the value to 70. However, you can successfully specify and set CI to any value over 70.

`ConfInterval` is also used with `RestartLimit` to prevent `vcsonclientd` from restarting the resource indefinitely. The `vcsonclientd` process attempts to restart the resource on the same system according to the number set in `RestartLimit` within `ConfInterval` before giving up and failing over. However, if the resource remains online for the interval designated in `ConfInterval`, earlier attempts to restart are not counted against `RestartLimit`. Default is 600 seconds.

The `ConfInterval` attribute value can be overridden.

## FaultOnMonitorTimeouts

Indicates the number of consecutive monitor failures to be treated as a resource fault. A monitor attempt is considered a failure if it does not complete within the time specified by the `MonitorTimeout` attribute.

When a monitor fails as many times as the value specified by this attribute, the corresponding resource is brought down by calling the `clean` entry point. The resource is then marked `FAULTED`, or it is restarted, depending on the value set in the `Restart Limit` attribute.

---

**Note:** This attribute applies only to online resources. If a resource is offline, no special action is taken during monitor failures.

---

When `FaultOnMonitorTimeouts` is set to 0, monitor failures are not considered indicative of a resource fault.

Default is 4. The `FaultOnMonitorTimeouts` attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## InfoInterval

Specifies the interval, in seconds, between successive invocations of the `info` entry point for a given resource. The default value of the `InfoInterval` attribute is 0, which specifies that the agent framework is not to schedule the `info` entry point periodically; the entry point can be invoked, however, by the user from the command line.

The `InfoInterval` attribute value can be overridden.

## InfoTimeout

A scalar integer specifying how long the agent framework allows for completion of the `info` entry point.

The default is 30 seconds. The value of the InfoTimeout attribute is internally capped at MonitorInterval / 2. The InfoTimeout attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## IntentionalOffline

Defines how VCS One reacts to a configured application being intentionally stopped outside of VCS One control.

Add this attribute for agents that support detection of an intentional offline outside of VCS One control.

Note that the intentional offline feature is available for agents registered as V51 or later. Use Script51Agent to use this feature.

The value 0 instructs the agent to register a fault and initiate the failover of a service group when the supported resource is taken offline outside of VCS One control. The default value for this attribute is 0.

The value 1 instructs VCS One to take the resource offline when the corresponding application is stopped outside of VCS One control. This attribute does not affect VCS One behavior on application failure. VCS One continues to fault resources if managed corresponding applications fail.

See “[About on-off, on-only, and persistent resources](#)” on page 16.

## Leveltwomonitorfrequency

The number of monitor cycles at which the agent framework initiates detailed monitoring. For example, if you set this attribute to 5, the agent framework initiates detailed monitoring every five monitor cycles.

## LogDbg

The LogDbg attribute indicates the debug severities enabled for the resource type or agent framework.

Debug severities used by agent functions are in the range of DBG\_1–DBG\_21. By default, LogDbg is an empty list, meaning that no debug messages are logged for a resource type. Users can modify this attribute for a given resource type, to

specify the debug severities that they want to enable, which would cause those debug messages to be printed to the log files.

The debug messages from the agent framework are logged with the following severities:

- `DBG_AGDEBUG`: Enables most debug logs.
- `DBG_AGINFO`: Enables debugging messages related to specific entry-point execution.
- `DBG_AGTRACE`: Enables verbose debug logging.

See [“Using debug messages”](#) on page 154.

See [“About logging agent messages”](#) on page 106 for information on the APIs available to log debug messages from agent entry points. These APIs expect a debug severity as a parameter, along with the message to be logged. You can choose different debug severities for messages to provide different logging levels for the agent. When you enable a particular severity in the `LogDbg` attribute, agent entry points log corresponding messages.

## LogFileSize

Sets the size of an agent log file. Value must be specified in bytes. Minimum is 65536 bytes (64KB). Maximum is 134217728 bytes (128MB). Default is 33554432 bytes (32MB). For example,

```
hatype -modify FileOnOff LogFileSize 2097152
```

Values specified less than the minimum acceptable value will be changed 65536 bytes. Values specified greater than the maximum acceptable value will be changed to 134217728 bytes. Therefore, out-of-range values displayed for the command:

```
hatype -display restype -attribute LogFileSize
```

will be those entered with the `-modify` option, not the actual values. The `LogFileSize` attribute value cannot be overridden.

## MinCompatAgentVersion

Specifies the lowest compatible version of the agent that is supported work with the current version of the type definition. The version specified in this attribute must be lower than or equal to the version specified in the `TypeDefinitionVersion` attribute. All versions of the agent binary starting from the version mentioned in this attribute upto the version given in the `TypeDefinitionVersion` attribute are assumed to be compatible with the type-definition.

## MonitorInterval

Duration (in seconds) between two consecutive monitor calls for an `ONLINE` resource or a resource in transition.

Default is 60 seconds. The `MonitorInterval` attribute value can be overridden.

## MonitorMethod

Internal attribute for future use. Do not modify.

## MonitorStatsParam

`MonitorStatsParam` is a type-level attribute, which stores the required parameter values for calculating monitor time statistics. For example:

```
static str MonitorStatsParam = { Frequency = 10, ExpectedValue =
 3000, ValueThreshold = 100, AvgThreshold = 40 }
```

- *Frequency*: Defines the number of monitor cycles after which the average monitor cycle time should be computed and sent to `vcstoneclientd`. The value of this key can be from 1 to 30. A value of 0 (zero) indicates that the average monitor time need not be computed. This is the default value for this key.
- *ExpectedValue*: The expected monitor time in milliseconds for all resources of this type. Default=100.
- *ValueThreshold*: The acceptable percentage difference between the expected monitor cycle time (*ExpectedValue*) and the actual monitor cycle time. Default=100.
- *AvgThreshold*: The acceptable percentage difference between the benchmark average and the moving average of monitor cycle times. Default=40.

The `MonitorStatsParam` attribute values can be overridden.

Refer to the *User's Guide*.

## MonitorTimeout

Maximum time (in seconds) within which the `monitor` entry point must complete or else be terminated. Default is 60 seconds. The `MonitorTimeout` attribute value can be overridden.

The determination of a suitable value for the `MonitorTimeout` attribute can be assisted by the use of the `MonitorStatsParam` attribute.

---

**Note:** This attribute is not used on Windows.

---

## NumThreads

NumThreads specifies the maximum number of service threads that an agent is allowed to create. Service threads are the threads in the agent that service resource commands, the main one being entry point processing. NumThreads does not control the number of threads used for other internal purposes.

Agents dynamically create service threads depending on the number of resources that the agent has to manage. Until the number of resources is less than the NumThreads value, the addition of a new resource will make the agent create an additional service thread. Also, if the number of resources falls below the NumThreads value as a result of deletion of resources, the agent will correspondingly delete service threads. Since an agent for a type will be started by VCS One `vcsonclientd` process only if there is at least one resource for that type in the configuration, an agent will always have at least 1 service thread. Setting NumThreads to 1 will thus prevent any additional service threads from being created even if more resources are added.

The maximum value that can be set for NumThreads is 30.

Default is 10. The NumThreads attribute *cannot* be overridden.

## OfflineMonitorInterval

The duration (in seconds) between two consecutive monitor calls for an `OFFLINE` resource. If set to 0, `OFFLINE` resources are not monitored.

Default is 0 seconds. The OfflineMonitorInterval attribute value can be overridden.

---

**Note:** Since the default value of this attribute is 0, concurrency violations are not detected. If an application that is supposed to be offline on a node is brought online outside of VCS One control, the application continues to run since VCS One cannot detect this state change. As mentioned, to avoid this, one can set OfflineMonitorInterval to a non-zero value (apart from overriding it for a specific resource).

---

## OfflineTimeout

Maximum time (in seconds) within which the `offline` entry point must complete or else be terminated.

Default is 300 seconds. The OfflineTimeout attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## OnlineRetryLimit

Number of times to retry `online` if the attempt to bring a resource online is unsuccessful. This attribute is meaningful only if `clean` is implemented.

Default is 0. The `OnlineRetryLimit` attribute value can be overridden.

## OnlineTimeout

Maximum time (in seconds) within which the `online` entry point must complete or else be terminated.

Default is 300 seconds. The `OnlineTimeout` attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## OnlineWaitLimit

Number of monitor intervals to wait after completing the online procedure, and before the resource is brought online. If the resource is not brought online after the designated monitor intervals, the online attempt is considered ineffective. This attribute is meaningful only if the `clean` entry point is implemented.

If `clean` is not implemented, the agent continues to periodically run `monitor` until the resource is brought online.

If `clean` is implemented, when the agent reaches the maximum number of monitor intervals it assumes that the online procedure was ineffective and runs `clean`. The agent then notifies `vcsonclientd` that the online attempt failed, or retries the procedure, depending on whether or not the `OnlineRetryLimit` is reached.

Default is 2. The `OnlineWaitLimit` attribute value can be overridden.

## OpenTimeout

Maximum time (in seconds) within which the `open` entry point must complete or else be terminated.

Default is 60 seconds. The `OpenTimeout` attribute value can be overridden.

---

**Note:** This attribute is not used on Windows.

---

## Operations

Indicates the valid operations for the resources of the type. The values are OnOff (can be brought online and taken offline), OnOnly (can be online only), and None (cannot be brought online or taken offline).

Default is OnOff. The Operations attribute value *cannot* be overridden.

## Platform

Specifies the platform for which this type is defined. A type could be defined for multiple platforms but the type's definition might be different for each platform. Hence the Platform attribute can be used to find out which platform the current definition is for.

The default value of this attribute, when a new type is added to a running cluster, is the platform specified in the DefaultPlatform cluster level attribute. If DefaultPlatform does not specify any platform, you must specify the platform when adding a new type using the “-platform” switch with the “hatype -add” command.

The Platform attribute *cannot* be overridden.

## RegList

RegList is a type level keylist attribute that can be used to store, or register, a list of certain resource level attributes. The agent calls the `attr_changed` entry point for a resource when the value of an attribute listed in RegList is modified. The RegList attribute is useful where a change in the values of important attributes require specific actions that can be executed from the `attr_changed` entry point.

By default, the attribute RegList is not included in a resource's type definition, but it can be added using either of the two methods shown below.

Assume the RegList attribute is added to the FileOnOff resource type definition and its value is defined as `PathName`. Thereafter, when the value of the `PathName` attribute for a FileOnOff resource is modified, the `attr_changed` entry point is called.

- Method one is to modify the types definition file (`types.xml` for example) to include the RegList attribute. Add a line in the definition of a resource type that resembles:

```
<attribute name="RegList" type="str" dimension="keylist">
 <default>
 <val key="attribute1_name"></val>
 <val key="attribute2_name"></val>
 </default>
</attribute>
```

For example, if the type definition is for the FileOnOff resource and the name of the attribute to register is PathName, the modified type definition would resemble:

```
<agent name="FileOnOff" platform="solaris/sparc" version="1.0">
.
.
.
 <attributes class="static">
 <attribute name="ArgList" type="str" dimension="vector">
 <default>
 <val>"PathName"</val>
 </default>
 <no_override>1</no_override>
 . . .
 </attribute>
 <attribute name="RegList" type="str" dimension="keylist">
 <default>
 <val key="PathName"></val>
 </default>
 </attribute>
 </attributes>
 <attributes>
 <attribute name="PathName" type="str" dimension="scalar">
 <important>1</important>
 <must_configure>1</must_configure>
 <unique>1</unique>
 <display_name>"PathName"</display_name>
 </attribute>
 </attributes>
</agent>
```

- Method two is to use the `haattr` command to add the RegList attribute to a resource type definition and then modify the value of the type's RegList attribute using the `hatype` command; the commands are:

```
haattr -add -static resource_type RegList -keylist -platform
platform_to_add_RegList_on
hatype -modify resource_type RegList attribute_name
 -platform platform_name
```

For example:

```
haattr -add -static FileOnOff RegList -keylist -platform
linux
hatype -modify FileOnOff RegList PathName -platform linux
```

The RegList attribute *cannot* be overridden.

## ResFaultPolicy

The `ResFaultPolicy` resource attribute controls how resources are to be handled in the event they become faulted. Each resource within a service group can have a policy. By default, VCS One assigns resources the `FaultPropagateAll` policy, whereby a faulting resource causes the entire service group and all resources to fault.

Possible values (or policies) for the resource attribute `ResFaultPolicy` include:

- **FaultNone:**  
The agent does not call the `clean` entry point.  
VCS One marks the resource as in the `ADMIN_WAIT` state. When the fault is removed and the `ADMIN_WAIT` state is cleared, VCS One fails over the service group.
- **FaultHold:**  
The agent calls the `clean` entry point.  
VCS One attempts to restart the resource according to the value of the `RestartLimit` resource attribute. VCS One does not propagate the fault to other resources in the service group.
- **FaultPropagateAll (default):**  
The agent calls the `clean` entry point.  
Before taking the parent group offline, VCS One attempts to restart the resource based on the value of the resource attribute `RestartLimit` before declaring the resource as faulted.  
Before taking the service group offline, VCS One attempts to restart the group based on the value of the resource attribute `OnlineRetryLimit`.  
VCS One takes the entire service group offline and attempts to fail over the service group, based on the value of the group attribute `GrpFaultPolicy`.

## RestartLimit

Affects how the agent responds to a resource fault.

Refer also to:

[“ToleranceLimit”](#) on page 180.

A non-zero value for `RestartLimit` causes the invocation of the `online` entry point instead of the failover of the service group to another system. `vcsonclientd` process attempts to restart the resource according to the number set in `RestartLimit` before it gives up and attempts failover. However, if the resource remains online for the interval designated in `ConfInterval`, earlier attempts to restart are not counted against `RestartLimit`.

---

**Note:** The agent will not restart a faulted resource if the `clean` entry point is not implemented. Therefore, the value of the `RestartLimit` attribute applies only if `clean` is implemented.

---

Default is 0. The `RestartLimit` attribute value can be overridden.

## ScriptClass

Indicates the scheduling class of the script processes (for example, online) created by the agent.

Default is “TS”.

## ScriptPriority

Indicates the priority of the script processes created by the agent.

Default is 0.

## SupportedActions

The SupportedActions (string-keylist) attribute lists all possible actions defined for an agent, including those defined by the agent developer. The `vcsonclientd` process validates the `action_token` value specified in the `hares -action resource action_token` command against the SupportedActions attribute. For example, if `action_token` is not present in SupportedActions, `vcsonclientd` will not allow the command to go through. It is the responsibility of the agent developer to initialize the SupportedActions attribute in the resource type definition and update the definition for each new action added to the `action` entry point code or script. This attribute serves as a reference for users of the command line or the graphical user interface.

See [“About the action entry point”](#) on page 33.

An example definition of a resource type in an VCS One

`ResourceTypeTypes.xml` file may resemble:

```
.
.
<agent name="DBResource" version = 1.0 bundled="False">
.
.
<attributes>
 <ArgList type="str" dimension="vector">
 <attr_description>" "</attr_description>
 <default>
 <val>Sid</val>
 <val>Owner</val>
 <val>Home</val>
 <val>User</val>
 <val>Pword</val>
 <val>StartOpt</val>
 <val>ShutOpt</val>
 </default>
 </ArgList>
.
.
 <SupportedActions type="str" dimension="keylist">
 <default>
 <val key="VRTS_GetRunningServices"/>
 <val key="DBRestrict"/>
 <val key="DBUndoRestrict"/>
 <val key="DBSuspend"/>
 <val key="DBResume"/>
 </default>
 </SupportedActions>
</attributes>
```

## ToleranceLimit

A non-zero `ToleranceLimit` allows the `monitor` entry point to return `OFFLINE` several times before the `ONLINE` resource is declared `FAULTED`. If the `monitor` entry point reports `OFFLINE` more times than the number set in `ToleranceLimit`, the resource is declared `FAULTED`. However, if the resource remains online for the interval designated in `ConfInterval`, any earlier reports of `OFFLINE` are not counted against `ToleranceLimit`. Default is 0. The `ToleranceLimit` attribute value can be overridden.

## TypeDefinitionVersion

Specifies the version of the type-definition for the custom type. If the agent code is modified or if the type-definition is changed in any way, then this attribute needs to be updated to a newer version. This attribute allows users to keep a track of the version of the type-definition for the type.

# State transition diagrams

- [State transitions](#)
- [State transitions for quiesce and unquiesce](#)
- [State transitions with respect to ResFaultPolicy attribute](#)

## State transitions

This section describes state transitions for:


- Opening a resource
- Resource in a steady state
- Bringing a resource online
- Taking a resource offline
- Resource fault (without automatic restart)
- Resource fault (with automatic restart)
- Monitoring of persistent resources
- Closing a resource

In addition, a set of diagrams describe the state transitions of resources with respect to the `ResFaultPolicy` resource attribute.

See “[State transitions with respect to ResFaultPolicy attribute](#)” on page 201.

The states shown in these diagrams are associated with each resource by the agent framework. These states are used only within the agent framework and are independent of the IState resource attribute values indicated by the engine. The agent writes resource state transition information into the agent log file when the `LogDbg` parameter, a static resource type attribute, is set to the value `DBG_AGINFO`. Agent developers can make use of this information when debugging agents.


### Opening a resource



When the agent starts up, each resource starts with the initial state of Detached. In the Detached state (`Enabled=0`), the agent rejects all commands to bring a resource online or take it offline.


When the resource is enabled (`Enabled=1`), the `open` entry point is invoked. Periodic Monitoring begins after `open` times out or succeeds. Depending on the return value of `monitor`, the resource transitions to either the Online or the Offline state. In the unlikely event that `monitor` times out or returns unknown, the resource stays in a Probing state.

## Resource in steady state



When resources are in a steady state of Online or Offline, they are monitored at regular intervals. The intervals are specified by the `MonitorInterval` parameter for a resource in the Online state and by the `OfflineMonitorInterval` parameter for a resource in the Offline state. An Online resource that is unexpectedly detected as Offline is considered to be faulted. Refer to diagrams describing faulted resources.

### Bringing a resource online: ResFaultPolicy = FaultPropagateAll




When the agent receives a request from the engine to bring the resource online, the resource enters the Going Online state, where the *online* entry point is invoked. If *online* completes successfully, the resource enters the Going Online Waiting state where it waits for the next monitor cycle.

If *monitor* returns a status of online, the resource moves to the Online state.

If, however, the *monitor* times out, or returns a status of “not Online” (that is, unknown or offline), the agent returns the resource to the Going Online Waiting state and waits for the next monitor cycle.

When the OnlineWaitLimit is reached, the *clean* entry point is invoked.

- If *clean* times out or fails, the resource again returns to the Going Online Waiting state and waits for the next monitor cycle. The agent again invokes the *clean* entry point if the monitor reports a status of “not Online.”
- If *clean* succeeds with the OnlineRetryLimit reached, and the subsequent monitor reports offline status, the resource transitions to the offline state and is marked FAULTED.
- If *clean* succeeds and the ORL is not reached, the resource transitions to the Going Online state where the *online* entry point is retried.



Upon receiving a request from the engine to take a resource offline, the agent places the resource in a Going Offline state and invokes the *offline* entry point.


If *offline* succeeds, the resource enters the Going Offline Waiting state where it waits for the next monitor.

If *monitor* returns a status of offline, the resource is marked Offline.

If the monitor times out or return a status “not offline,” the agent invokes the *clean* entry point. Also, if, in the Going Offline state, the *offline* entry point times out, the agent invokes *clean* entry point.

- If *clean* fails or times out, the resource is placed in the Going Offline Waiting state and monitored. If *monitor* reports “not offline,” the agent invokes the *clean* entry point, where the sequence of events repeats.
- If *clean* returns success, the resource is placed in the Going Offline Waiting state and monitored. If *monitor* times out or reports “not offline,” the resource returns to the GoingOfflineWaiting state. The UNABLE\_TO\_OFFLINE flag is sent to engine.

### Resource fault without automatic restart



This diagram describes the activity that occurs when a resource faults and the `RestartLimit` is reached. When the `monitor` entry point times out successively and `FaultOnMonitorTimeout` is reached, or `monitor` returns offline and the `ToleranceLimit` is reached, the agent invokes the `clean` entry point.

---

**Note:** For the Windows platform, the `FaultOnMonitorTimeout` attribute has no significance. Instead, the `monitor` entry point is allowed to continue to completion.


---

If `clean` fails, or if it times out, the agent places the resource in the Online state as if no fault occurred.

If `clean` succeeds, the resource is placed in the Going Offline Waiting state, where the agent waits for the next monitor.

- If `monitor` reports Online, the resource is placed back Online as if no fault occurred. If `monitor` reports Offline, the resource is placed in an Offline state and marked as FAULTED.
- If `monitor` reports unknown or times out, the agent places the resource back into the Going Offline Waiting state, and sets the `UNABLE_TO_OFFLINE` flag in the engine.

### Resource fault with automatic restart



This diagram describes the activity that occurs when a resource faults and the `RestartLimit` is not reached. When the `monitor` entry point times out successively and `FaultOnMonitorTimeout` is reached, or `monitor` returns offline and the `ToleranceLimit` is reached, the agent invokes the `clean` entry point.

---


**Note:** For the Windows platform, the `FaultOnMonitorTimeout` attribute has no significance. Instead, the `monitor` entry point is allowed to continue to completion.

---

- If `clean` succeeds, the resource is placed in the `Going Online` state and the `online` entry point is invoked to restart the resource; refer to the diagram, “Bringing a resource online.”
- If `clean` fails or times out, the agent places the resource in the `Online` state as if no fault occurred.


Refer to the diagram “Resource fault without automatic restart,” for a discussion of activity when a resource faults and the `RestartLimit` is reached.

### Monitoring of persistent resources



For a persistent resource in the Online state, if *monitor* returns a status of offline and the `ToleranceLimit` is not reached, the resource stays in an Online state. If *monitor* returns offline and the `ToleranceLimit` is reached, the resource is placed in an Offline state and noted as FAULTED. If *monitor* returns “not offline,” the resource stays in an Online state.

Likewise, for a persistent resource in an Offline state, if *monitor* returns offline, the resource remains in an Offline state. If *monitor* returns a status of online, the resource is placed in an Online state.




When the resource is disabled (Enabled=0), the agent stops Periodic Monitoring and the *close* entry point is invoked. When the *close* entry point succeeds or times out, the resource is placed in the Detached state.

## State transitions for quiesce and unquiesce

An agent moves resources to a quiesced state when vcsonclientd loses connection with the policy master. When the connection is re-established, the agent unquiesces the resources. In a quiesced state, none of the resource commands work. The agent unquiesces resources only after vcsonclientd instructs the agent to unquiesce its resources.

### Quiescing and unquiescing of a resource in Detached state




Re-init: clear all previous resource information except snapshot information.

Quiescing and unquiescing of a resource in Cleaning state




Legend
[ <i>state</i> ] => represents the previous agfw- istate of the resource
d_w_w_q = done_waiting_while_quiesced flag at resource-object level

Quiescing and unquiescing of a resource in Closing state




Legend
[ <i>state</i> ] => represents the previous agfw- istate of the resource
d_w_w_q = done_waiting_while_quiesced flag at resource-object level


### Quiescing and unquiescing of a resource in GoingOffline state



Quiescing and unquiescing of a resource in GoingOnline state




### Quiescing and unquiescing of a resource in GoingOnlineWaiting state




Legend
[ <istate> ] => represents the previous agfw-istate of the resource
d_w_w_q = done_waiting_while_quiesced flag at resource-object level

Quiescing and unquiescing of a resource in Monitoring state




### Quiescing and unquiescing of a resource in Offline state




- 1) Insert monitor to check if resource came online outside of VCSONe control
- 2) Periodic offline monitoring is already on, if OfflineMonitorInterval is non-zero; so no need to setup periodic monitoring.

Legend
[ <istate> ] => represents the previous agfw-istate of the resource
d_w_w_q = done_waiting_while_quiesced flag at resource-object level


Quiescing and unquiescing of a resource in Online state




### Quiescing and unquiescing of a resource in Opening state



### Quiescing and unquiescing of a resource in Probing state



Quiescing and unquiescing of a resource in Quiesced state



## State transitions with respect to ResFaultPolicy attribute


This section show state transition diagrams with respect to the ResFaultPolicy attribute.

By default, the ResFaultPolicy attribute is set to FaultPropagateAll, in which case the agent calls the clean entry point. Based on the value of the resource type level attribute RestartLimit, the agent attempts to restart the resource before declaring it as FAULTED. In general, when ResFaultPolicy attribute is set to FaultHold, the behavior is the same as when it is set to FaultPropagateAll. Further actions depend on the group level attribute GrpFaultPolicy.


The next set of diagrams illustrate the following state transitions:

- Bringing a resource online, ResFaultPolicy set to FaultNone
- Taking a resource offline, ResFaultPolicy set to FaultNone
- Resource fault, ResFaultPolicy set to FaultPropagateAll or FaultHold
- Resource fault (unexpected offline), ResFaultPolicy set to FaultNone
- Resource fault (monitor hung), ResFaultPolicy set to FaultPropagateAll
- Resource fault (monitor hung), ResFaultPolicy set to FaultNone


Bringing a resource online: ResFaultPolicy = FaultNone




### Taking a resource offline; ResFaultPolicy = FaultNone



Resource fault: ResFaultPolicy = FaultPropagateAll or FaultHold



### Resource fault (unexpected offline): ResFaultPolicy = FaultNone




PM - Periodic Monitoring  
TL - Tolerance Limit



### Resource fault (monitor hung): ResFaultPolicy = FaultNone

MSG\_RES\_CLEARADM\_WAIT  
or message to set  
ResFaultPolicy to  
FaultPropagateAll or  
FaultHold

MSG\_RES\_CLRADMIN\_WAIT\_FAULT /  
reset ADMIN\_WAIT;  
faulted = 1





## Using VCS agents with VCS One

- [About using VCS agents with VCS One](#)
- [Using pre-5.0 VCS agents with VCS One and registering them as V51](#)

## About using VCS agents with VCS One

The V51 agent framework used by VCS One supports the use of pre-50 VCS agents

## Using pre-5.0 VCS agents with VCS One and registering them as V51

When you use pre-5.0 agents with VCS One, you may register them as V51 agents after making necessary modifications. Making this conversion affords you advantages, which include:

- You can use different versions of an agent on different systems in the cluster.
- You can make changes to the resource type definition used on some systems without affecting how older versions of the agents function

### Outline of steps to change V40 agents V51

- Change the way attributes and their values are passed to the entry points from the V40 format to V51 name-value tuple format.
- Include /opt/VRTSvcs/lib in path for Perl and shell to source them.
- Set necessary environment variables.

### Example script in V40 and V51

Note the following comparison.

#### V40

```
ResName=$1
Attr1=$2
Attr2=$3
VCSHOME="{VCSONE_HOME:-/opt/VRTSvcsone}"
. $VCSHOME/bin/ag_i18n_inc.sh;
VCSAG_SET_ENVS $ResName;
```

#### V51

```
ResName=$1; shift;
.."./ag_i18n_inc.sh";
VCSAG_SET_ENVS $ResName;

VCSAG_GET_ATTR_VALUE "Attr1" -1 1 "$@";
attr1_value=${VCSAG_ATTR_VALUE};
VCSAG_GET_ATTR_VALUE "Attr2" -1 1 "$@";
```

```
attr2_value=${VCSAG_ATTR_VALUE};
```

## Sourcing ag\_i18n\_inc modules in script entry points

In entry points, you need to source the ag\_i18n\_inc modules. The following examples assume that the agent is installed in the directory */opt/VRTSvcs/bin/type*.

For entry points in Perl:

```
...
$ResName = shift;
use ag_i18n_inc;
VCSAG_SET_ENVS ($ResName);
...
```

For entry points in Shell:

```
...
ResName = $1; shift;
. "../ag_i18n_inc.sh";
VCSAG_SET_ENVS $ ResName;
```



# Index

## A

- action entry point
  - script entry point 98
  - supported actions 179
- agent framework
  - library, C++ 37
  - logging APIs 108
  - multithreaded 46
  - Windows primitives 85
- agent messages
  - for Windows 85
  - formatting 107
  - normal in VCSAG\_LOG\_MSG 121
- AgentClass parameter 164
- AgentPriority parameter 165
- AgentReplyTimeout parameter 165
- Appendix 106
- ArgList parameter 165
- ArgList reference attributes 165
- association attribute dimension 20
- attr\_changed entry point 34
  - C++ syntax 60
  - script syntax 98
- AttrChangedTimeout parameter 166
- attribute dimensions
  - association 20
  - keylist 19
  - scalar 18
  - vector 19
- attribute types
  - boolean 18
  - integer 18
  - string 18
- attributes
  - local and global 20

## B

- boolean attribute type 18

## C

- category ID for messages 118
- clean entry point 31
  - C++ syntax 57
  - enum types 31
  - script syntax 98
- CleanTimeout parameter 166
- close entry point 35
  - C++ syntax 62
  - script syntax 99
- CloseTimeout parameter 167
- configuration language
  - local and global attributes 20
- ConfInterval parameter 168

## D

- debug message severity level 110
- debug messages
  - C++ entry points 110
  - Perl script entry points 122
  - Shell script entry points 122

## E

- entry points
  - attr\_changed 34
  - clean 31
  - close 35
  - definition 16
  - info 28
  - monitor 27
  - offline 30
  - online 30
  - open 34
  - sample structure 38
  - shutdown 35
- enum types for clean
  - VCSAgCleanMonitorHung 32
  - VCSAgCleanOfflineHung 31
  - VCSAgCleanOfflineIneffective 31
  - VCSAgCleanOnlineHung 31

VCSAgCleanOnlineIneffective 32  
 VCSAgCleanUnexpectedOffline 32

## F

FaultOnMonitorTimeouts parameter 169  
 formatting agent messages 107

## G

global attributes 20

## I

info entry point 28  
     C++ syntax 51  
     script example 99  
 InfoTimeout parameter 170  
 initializing functions with VCSAG\_LOG\_INIT 112  
 integer attribute type 18  
 intentional offline 16

## K

keylist attribute dimension 19

## L

local attributes 20  
 log category 113  
 LogDbg parameter 170  
 LogFileSize parameter 171  
 logging APIs  
     C++ 108  
     script entry points 118

## M

message text format 107, 108  
 mnemonic message field 107  
 monitor entry point 27  
     C++ syntax 50, 51  
     script syntax 97  
 MonitorLevel parameter 172  
 MonitorStatsParam parameter 172  
 MonitorTimeout parameter 172

## O

offline entry point 30  
     C++ syntax 56  
     script syntax 97

OfflineMonitorInterval parameter 173  
 OfflineTimeout parameter 173  
 online entry point 30  
     C++ syntax 55  
     script syntax 97  
 OnlineRetryLimit parameter 174  
 OnlineTimeout parameter 174  
 OnlineWaitLimit parameter 174  
 OnOff resource type 16  
 OnOnly resource type 16  
 open entry point 34  
     C++ syntax 62  
     script syntax 99  
 OpenTimeout parameter 174  
 Operations parameter 174

## P

parameters  
     AgentClass 164  
     AgentPriority 165  
     AgentReplyTimeout 165  
     ArgList 165  
     AttrChangedTimeout 166  
     CleanTimeout 166  
     CloseTimeout 167  
     ConfInterval 168  
     FaultOnMonitorTimeouts 169  
     InfoTimeout 170  
     LogDbg 170  
     LogFileSize 171  
     MonitorLevel 172  
     MonitorStatsParam 172  
     MonitorTimeout 172  
     OfflineMonitorInterval 173  
     OfflineTimeout 173  
     OnlineRetryLimit 174  
     OnlineTimeout 174  
     OnlineWaitLimit 174  
     OpenTimeout 174  
     Operations 174  
     RegList 175  
     RestartLimit 177  
     ScriptClass 178  
     ScriptPriority 178  
     ToleranceLimit 180  
 persistent resource type 16  
 primitives  
     definition 65  
     logging primitives for Windows 85

- VCSAgAllocMem 88
- VCSAgDelEncodedArgList 87
- VCSAGDeleteMem 88
- VCSAgGetCookie 69
- VCSAgGetEncodedArglist 87
- VCSAgRegister 67
- VCSAgRegisterEPStruct 65
- VCSAgSetCookie 65
- VCSAgSetCookie2 65
- VCSAgSnprintf 70
- VCSAgSnprintfW 87
- VCSAgStrlcat 70
- VCSAgStrlcatW 87
- VCSAgStrlcpyW 87
- VCSAgUnregister 68

## Q

- quiesce, state transitions 191

## R

- RegList parameter 175
- resource
  - closing (state transition diagram) 190
  - fault (state transition diagram) 187, 188
  - monitoring (state transition diagram) 189
  - offlining (state transition diagram) 186
  - onlining (state transition diagram) 185
  - OnOff type 16
  - OnOnly type 16
  - opening (state transition diagram) 183
  - persistent type 16
  - steady state (state transition diagram) 184
- RestartLimit parameter 177

## S

- scalar attribute dimension 18
- Script50Agent 141
- script-based logging functions 117
- ScriptClass parameter 178
- ScriptPriority parameter 178
- severity macros 111
- severity message field 107
- shutdown entry point 35
  - C++ syntax 64
  - script syntax 99
- state transition diagram
  - closing a resource 190

- monitoring persistent resources 189
- offlining a resource 186
- onlining a resource 185
- opening a resource 183
- resource fault with auto restart 188
- resource fault, no auto restart 187
- resource in steady state 184

- string attribute type 18

## T

- timestamp message field 107
- ToleranceLimit parameter 180

## U

- UCS-2 encoding (wchar)
  - enabling agents with 85
  - primitives for converting to 86
- UMI (unique message identifier) 107
- unquiesce, state transitions 191

## V

- VCSAG\_CONSOLE\_LOG\_MSG logging macro 108
- VCSAG\_LOG\_INIT initializing function 112
- VCSAG\_LOG\_MSG logging macro 108
- VCSAG\_LOG\_MSG script logging function 117
- VCSAG\_LOGDBG\_MSG logging macro 108
- VCSAG\_LOGDBG\_MSG script logging function 117
- VCSAG\_RES\_LOG\_MSG logging macro 108
- VCSAG\_SET\_ENVS script logging function 117
- VCSAgAllocMem primitive 88
- VCSAgDelEncodedArgList primitive 87
- VCSAgDeleteMem primitive 88
- VCSAgGetCookie primitive 69
- VCSAgGetEncodedArglist primitive 87
- VCSAgRegister primitive 67
- VCSAgRegisterEPStruct primitive 65
- VCSAgSetCookie primitive 65
- VCSAgSetCookie2 primitive 65
- VCSAgSnprintf primitive 70
- VCSAgSnprintfW primitive 87
- VCSAgStartup entry point, C++ syntax 49
- VCSAgStrlcat primitive 70
- VCSAgStrlcatW primitive 87
- VCSAgStrlcpyW primitive 87
- VCSAgUnregister primitive 68
- VCSAgValidateAndSetEntryPoint 77
- vector attribute dimension 19

