

Cluster Server Agent for IBM InfoSphere DataStage Installation and Configuration Guide

Linux, Solaris

7.0

VERITAS™

Veritas InfoScale™ Availability Agents

Last updated: 2020-01-09

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introducing the agent for IBM InfoSphere DataStage	6
	About the Cluster Server agent for IBM InfoSphere DataStage	6
	Supported software	7
	Features of the agent	7
	How the agent makes IBM InfoSphere DataStage highly available	7
	How the agent supports intelligent resource monitoring	8
	IBM InfoSphere DataStage agent functions	8
	Online	9
	Offline	9
	Monitor	10
	Clean	10
	imf_init	11
	imf_getnotification	11
	imf_register	11
Chapter 2	Installing, upgrading, and removing the agent for IBM InfoSphere DataStage	12
	Before you install the Cluster Server agent for IBM InfoSphere DataStage	12
	About the ACC library	13
	Installing the ACC library	13
	Installing the agent in a VCS environment	14
	Uninstalling the agent in a VCS environment	15
	Removing the ACC library	16
Chapter 3	Configuring the agent for IBM InfoSphere DataStage	17
	About configuring the Cluster Server agent for IBM InfoSphere DataStage	17
	Importing the agent types file	17
	InfoSphere agent attributes	18
	About the keys of the IMF attribute	22

	Attributes used in different resource configurations	23
	Executing a customized monitoring program	24
	Enabling and disabling intelligent resource monitoring manually	25
Chapter 4	Troubleshooting the agent for IBM InfoSphere	
	DataStage	27
	Using the correct software and operating system versions	27
	Meeting prerequisites	27
	Reviewing error log files	28
	Using InfoSphere log files	28
	Reviewing cluster log files	28
	Using trace level logging	28
	Troubleshooting the configuration for IMF	29
	Known issues	31
Appendix A	Sample Configurations	33
	About sample configurations for the agents for IBM InfoSphere	
	DataStage	33
	Sample agent type definition for InfoSphere	33
	Sample configuration files	34
	Sample service group configurations	36
	Sample resource configurations for InfoSphere	37

Introducing the agent for IBM InfoSphere DataStage

This chapter includes the following topics:

- [About the Cluster Server agent for IBM InfoSphere DataStage](#)
- [Supported software](#)
- [Features of the agent](#)
- [How the agent makes IBM InfoSphere DataStage highly available](#)
- [How the agent supports intelligent resource monitoring](#)
- [IBM InfoSphere DataStage agent functions](#)

About the Cluster Server agent for IBM InfoSphere DataStage

The Cluster Server (VCS) agent for IBM InfoSphere DataStage provides high availability for InfoSphere DataStage components in a cluster. The agent monitors specific InfoSphere DataStage components, such as DS Engine, ASB Agent, and ASB Logging Agent. The agent brings these components online, takes them offline and, in case of a failure, shuts the components down.

Note: In InfoSphere 11.5, the functionalities of the ASBLoggingAgent and the ASBNodeAgent components have been merged into ASBNodeAgent. Accordingly, Veritas has also merged the ASBLoggingAgent and the ASBNodeAgent components of the VCS agent for InfoSphere DataStage. The ASBLoggingAgent component is no longer applicable in the newer versions of InfoSphere.

Supported software

For information on the software versions that the Cluster Server agent for IBM InfoSphere DataStage supports, see the Veritas Services and Operations Readiness Tools (SORT) site: <https://sort.veritas.com/agents>.

Features of the agent

The following are the features of the Cluster Server agent for IBM InfoSphere DataStage:

- Support for validation of attributes that are based on the agent functions
The agent can validate attributes in each agent function before the actual data processing starts.
- Support for First Failure Data Capture (FFDC)
In case of a fault, the agent generates a huge volume of the debug logs that enable troubleshooting of the fault.
- Support for Fast First Level Monitor (FFLM)
The agent maintains PID files based on search patterns to expedite the monitoring process.
- Support for external user-supplied monitor utilities
The agent enables user-specified monitor utilities to be plugged in, in addition to the built-in monitoring logic. This enables administrators to completely customize the monitoring of the application.
- Support for intelligent resource monitoring and poll-based monitoring
The agent supports the Cluster Server Intelligent Monitoring Framework (IMF) feature. IMF allows the agent to register the resources to be monitored with the IMF notification module so as to receive immediate notification of resource state changes without having to periodically poll the resources.
- Delayed agent function
The agent manages the first monitor after online for slow initializing applications.

How the agent makes IBM InfoSphere DataStage highly available

The agent provides the following levels of application monitoring:

- Primary or Basic monitoring
This mode has Process check and Health check monitoring options. With the default Process check option, the agent verifies that the InfoSphere DataStage

component processes are present in the process table. Process check cannot detect whether processes are in the hung or stopped states.

- Secondary or Detail monitoring

In this mode, the agent runs a utility to verify the status of the InfoSphere DataStage component. The agent detects application failure if the monitoring routine reports an improper function of the InfoSphere DataStage component processes. When this application failure occurs, the InfoSphere DataStage component service group fails over to another node in the cluster.

In addition to these levels of application monitoring, the agent for IBM InfoSphere DataStage is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification.

Thus, the agent ensures high availability for InfoSphere DataStage components.

How the agent supports intelligent resource monitoring

With Intelligent Monitoring Framework (IMF), VCS supports intelligent resource monitoring in addition to the poll-based monitoring. Poll-based monitoring polls the resources periodically whereas intelligent monitoring performs asynchronous monitoring.

When an IMF-enabled agent starts up, the agent initializes the Asynchronous Monitoring Framework (AMF) kernel driver. After the resource is in a steady state, the agent registers with the AMF kernel driver, the details of the resource that are required to monitor the resource. For example, the agent for IBM InfoSphere DataStage registers the PIDs of the InfoSphere processes with the AMF kernel driver. The agent's `imf_getnotification` function waits for any resource state changes. When the AMF kernel driver module notifies the `imf_getnotification` function about a resource state change, the agent framework runs the monitor agent function to ascertain the state of that resource. The agent notifies the state change to VCS, which then takes appropriate action.

You can enable or disable the intelligent resource monitoring functionality of the VCS agent for Cluster Server manually.

See [“Enabling and disabling intelligent resource monitoring manually”](#) on page 25.

For more information, see the *Cluster Server Administrator's Guide*.

IBM InfoSphere DataStage agent functions

The operations or functions that the Cluster Server Agent for InfoSphere can perform are described as follows:

Online

The online function performs the following tasks:

- Verifies that the required attributes are set correctly.
- Verifies that the InfoSphere DataStage component is not already online. If the component is online, the online operation exits immediately.
- If any InfoSphere DataStage component processes remain, the operation kills these processes using the user name associated with the specific resource.
- Depending on the InfoSphere DataStage component that is configured, the agent attempts to start the InfoSphere DataStage component with the command:

```
For DSEngine      InfoSphereInstanceDir/Server/DSEngine/bin/uv\  
                  -admin start
```

```
ASBNodeAgent     InfoSphereInstanceDir/ASBNode/bin/NodeAgents.sh\  
                  start
```

```
ASBLoggingAgent  InfoSphereInstanceDir/ASBNode/bin/NodeAgents.sh\  
                  startLogging
```

The command always gets executed in the context of User.

- Checks if the server has started up completely.
- Gives the control back to HAD.

Offline

The offline function performs the following tasks:

- Verifies that the required attributes are set correctly.
- Verifies that the InfoSphere DataStage component is not offline.
If the instance is already offline, the operation verifies if any processes belonging to this InfoSphere resource exist.
- Depending on the InfoSphere DataStage component that is configured, the agent attempts to stop the InfoSphere DataStage component with the command:

```
For DSEngine      InfoSphereInstanceDir/Server/DSEngine/bin/uv\  
                  -admin stop
```

```
ASBNodeAgent InfoSphereInstanceDir/ASBNode/bin/NodeAgents.sh\  
stopAgent
```

```
ASBLoggingAgent InfoSphereInstanceDir/ASBNode/bin/NodeAgents.sh\  
stop
```

The command always gets executed in the context of User.

- Kills any existing processes that belong to this InfoSphere DataStage component.
- Gives the control back to HAD.

Monitor

The monitor function monitors the states of the InfoSphere servers on all nodes within the cluster. The operation performs the following tasks:

- The monitor function conducts a first-level check to determine that the InfoSphere server processes are running on the system in the cluster. If the first-level check does not find these processes running on the node, the check exits immediately and reports the instance as OFFLINE.

The agent for InfoSphere also supports Intelligent Monitoring Framework (IMF) in the first-level check. IMF enables intelligent resource monitoring. See [“How the agent supports intelligent resource monitoring”](#) on page 8. You can use the MonitorFreq key of the IMF attribute to specify the frequency at which the agent invokes the monitor function.

- If in-depth monitoring is enabled, the monitor operation conducts a second-level check. During this check, the agent attempts to see if the InfoSphere DataStage component is listening on the port of a particular host by using the socket method.
- Depending upon the value of the MonitorProgram attribute, the monitor operation can perform a customized check using a user-supplied monitoring utility.

Clean

In case of a failure or after an unsuccessful attempt to bring an InfoSphere DataStage component online or take an InfoSphere DataStage component offline, the clean operation performs the following tasks:

- Depending on the InfoSphere DataStage component that is configured, the agent attempts to gracefully shut down the InfoSphere DataStage component with the command:

```
For DSEngine      InfoSphereInstanceDir/Server/DSEngine/bin/uv\  
                  -admin stop  
  
ASBNodeAgent     InfoSphereInstanceDir/ASBNode/bin/NodeAgents.sh\  
                  stopAgent  
  
ASBLoggingAgent  InfoSphereInstanceDir/ASBNode/bin/NodeAgents.sh\  
                  stop
```

The command always gets executed in the context of User.

- Kills any remaining processes pertaining to this InfoSphere instance.
- Gives the control back to HAD.

imf_init

This function initializes the IBM InfoSphere DataStage agent to interface with the AMF kernel driver, which is the IMF notification module for the agent for IBM InfoSphere DataStage. This function runs when the agent starts up.

imf_getnotification

This function gets notifications about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.

imf_register

This function registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into a steady state—online or offline.

Installing, upgrading, and removing the agent for IBM InfoSphere DataStage

This chapter includes the following topics:

- [Before you install the Cluster Server agent for IBM InfoSphere DataStage](#)
- [About the ACC library](#)
- [Installing the ACC library](#)
- [Installing the agent in a VCS environment](#)
- [Uninstalling the agent in a VCS environment](#)
- [Removing the ACC library](#)

Before you install the Cluster Server agent for IBM InfoSphere DataStage

You must install the Cluster Server agent for IBM InfoSphere DataStage on all the systems that will host InfoSphere service groups.

Before you install the agent for IBM InfoSphere DataStage, ensure that the following prerequisites are met.

- Install and configure Cluster Server.
For more information on installing and configuring Cluster Server, refer to the Cluster Server installation and configuration guides.
- Install the latest version of ACC Library.

To install or update the ACC Library package, locate the library and related documentation in the Agent Pack tarball.
See “[About the ACC library](#)” on page 13.

About the ACC library

The operations of a Cluster Server agent depend on a set of Perl modules known as the ACC library. The library must be installed on each system in the cluster that runs the agent. The ACC library contains common, reusable functions that perform tasks, such as process identification, logging, and system calls.

Instructions to install or remove the ACC library on a single system in the cluster are given in the following sections. The instructions assume that the ACCLib tar file has already been extracted.

Note: The LogDbg attribute should be used to enable debug logs for the ACCLib-based agents when the ACCLib version is 6.2.0.0 or later and VCS version is 6.2 or later.

Installing the ACC library

Install the ACC library on each system in the cluster that runs an agent that depends on the ACC library.

To install the ACC library

- 1 Log in as a superuser.
- 2 Download ACC Library.
You can download either the complete Agent Pack tar file or the individual ACCLib tar file from the Veritas Services and Operations Readiness Tools (SORT) site (<https://sort.veritas.com/agents>).
- 3 If you downloaded the complete Agent Pack tar file, navigate to the directory containing the package for the platform running in your environment.

```
Solaris      cd1/solaris/dist_arch/vcs/application/acc_library/version_library/pkg
```

- 4 If you downloaded the individual ACCLib tar file, navigate to the pkgs directory (for AIX and Solaris), or rpms directory (for Linux).
- 5 Install the package. Enter **Yes**, if asked to confirm overwriting of files in the existing package.

```
Solaris      # pkgadd -d VRTSacclib.pkg
```

Note: The LogDbg attribute should be used to enable debug logs for the ACCLib-based agents when the ACCLib version is 6.2.0.0 or later and VCS version is 6.2 or later.

Installing the agent in a VCS environment

Install the agent for IBM InfoSphere DataStage on each node in the cluster.

To install the agent in a VCS environment

- 1 Download the agent from the Veritas Services and Operations Readiness Tools (SORT) site: <https://sort.veritas.com/agents>.

You can download either the complete Agent Pack tar file or an individual agent tar file.

- 2 Uncompress the file to a temporary location, say /tmp.
- 3 If you downloaded the complete Agent Pack tar file, navigate to the directory containing the package for the platform running in your environment.

```
Solaris      cd1/solaris/dist_arch/vcs/application/infosphere_agent/  
             vcs_version/version_agent/pkgs
```

If you downloaded the individual agent tar file, navigate to the pkgs directory (for AIX and Solaris), or rpms directory (for Linux).

- 4 Log in as a superuser.
- 5 Install the package.

```
Solaris      # pkgadd -d . VRTSinfosphere
```

After installing the agent package, you must import the agent type configuration file.

See [“Importing the agent types file”](#) on page 17.

Uninstalling the agent in a VCS environment

You must uninstall the agent for IBM InfoSphere DataStage from a cluster while the cluster is active.

To uninstall the agent in a VCS environment

- 1 Log in as a superuser.
- 2 Set the cluster configuration mode to read/write by running the following command from any node in the cluster:

```
# haconf -makerw
```

- 3 Remove all InfoSphere resources from the cluster. Run the following command to verify that all resources have been removed:

```
# hares -list Type=InfoSphere
```

- 4 Remove the agent type from the cluster configuration by running the following command from any node in the cluster:

```
# hatype -delete InfoSphere
```

Removing the agent's type file from the cluster removes the include statement for the agent from the `main.cf` file, but the agent's type file is not removed from the cluster configuration directory. You can remove the agent's type file later from the cluster configuration directory.

- 5 Save these changes. Then set the cluster configuration mode to read-only by running the following command from any node in the cluster:

```
# haconf -dump -makero
```

- 6 Use the platform's native software management program to remove the agent for IBM InfoSphere DataStage from each node in the cluster.

Run the following command to uninstall the agent:

```
Linux      # rpm -e VRTSinfosphere
```

```
Solaris    # pkgrm VRTSinfosphere
```

Removing the ACC library

Perform the following steps to remove the ACC library.

To remove the ACC library

- 1 Ensure that all agents that use ACC library are removed.
- 2 Run the following command to remove the ACC library package:

```
Solaris          # pkgrm VRTSacclib
```

Configuring the agent for IBM InfoSphere DataStage

This chapter includes the following topics:

- [About configuring the Cluster Server agent for IBM InfoSphere DataStage](#)
- [Importing the agent types file](#)
- [InfoSphere agent attributes](#)
- [Executing a customized monitoring program](#)
- [Enabling and disabling intelligent resource monitoring manually](#)

About configuring the Cluster Server agent for IBM InfoSphere DataStage

After installing the Cluster Server agent for IBM InfoSphere DataStage, you must import the agent type configuration file. After importing this file, review the attributes table that describes the resource type and its attributes, and then create and configure InfoSphere resources.

To view the sample agent type definition and service groups configuration:

See [“About sample configurations for the agents for IBM InfoSphere DataStage”](#) on page 33.

Importing the agent types file

To use the agent for IBM InfoSphere DataStage, you must import the agent types file into the cluster.

You can import the agent types file using the Cluster Server graphical user interface or via the command line interface.

To import the agent types file using the VCS graphical user interface

- 1 Start the Cluster Manager and connect to the cluster on which the agent is installed.
- 2 Click **File > Import Types**.
- 3 In the Import Types dialog box, select the following file:

```
Linux /etc/VRTSagents/ha/conf/InfoSphere/InfoSphereTypes.cf
```

```
Solaris SPARC /etc/VRTSagents/ha/conf/InfoSphere/InfoSphereTypes51.cf
```

- 4 Click **Import**.
- 5 Save the VCS configuration.

The InfoSphere agent type is now imported to the VCS engine.

You can now create InfoSphere resources. For additional information about using the VCS GUI, refer to the *Cluster Server Administrator's Guide*.

To import the agent types file using the Cluster Server command line interface (CLI):

- 1 Log on to any one of the systems in the cluster as the superuser.
- 2 Run the following command:

```
# sh /etc/VRTSagents/ha/conf/InfoSphere/InfoSphereTypes.cmd
```

- 3 To verify that the agent types file is successfully imported to the VCS engine, run the following command:

```
# hatype -display InfoSphere
```

You can now create InfoSphere resources.

InfoSphere agent attributes

Refer to the required and optional attributes while configuring the agent for InfoSphere.

[Table 3-1](#) lists the required attributes for the InfoSphere agent.

Table 3-1 Required attributes

Attribute	Description
ResLogLevel	<p>Specifies the logging detail that the agent performs for the resource.</p> <p>The valid values are as follows:</p> <ul style="list-style-type: none"> ■ ERROR: Only logs error messages. ■ WARN: Logs error messages and warning messages. ■ INFO: Logs error messages, warning messages, and informational messages ■ TRACE: Logs error messages, warning messages, informational messages, and trace messages. TRACE is very verbose and should be used only during initial configuration or for troubleshooting and diagnostic operations. <p>Default Value: INFO</p> <p>Example: INFO</p>
Component	<p>Specifies the engine tier component for which the resource must be configured.</p> <p>The possible values are as follows:</p> <ul style="list-style-type: none"> ■ DSEngine ■ ASBNodeAgent ■ ASBLoggingAgent <p>Default Value: No default value</p> <p>Example: DSEngine</p>
User	<p>The OS user to run commands for starting engine tier components, such as DSEngine, ASBNode Agent and ASB Logging Agent.</p> <p>Default Value: No default value</p> <p>Example: root, dsadm</p>
InfoSphereInstanceDir	<p>The absolute path of the directory where the DataStage instance of the InfoSphere Information Server is installed.</p> <p>Default Value: No default value</p> <p>Example: /opt/IBM/InformationServer</p>

Table 3-1 Required attributes (*continued*)

Attribute	Description
InfoSphereInstance	<p>The unique identification tag (ITAG) or the instance name that was specified when the instance was created.</p> <p>You must configure this attribute only if the DSEngine InfoSphere DataStage component is configured.</p> <p>Default Value: No default value</p> <p>Example: ade</p>
EnvFile	<p>Full path of the file that the agent sources to set the environment before executing any DataStage server commands.</p> <p>Veritas recommends storing the file on shared disk. The following shell environments are supported: ksh, sh, and csh.</p> <p>You must configure this attribute only if the DSEngine InfoSphere DataStage component is configured.</p> <p>Default Value: No default value</p> <p>Example:</p> <pre data-bbox="602 857 1220 878">/opt/IBM/InformationServer/Server/DSEngine/dsenv</pre> <p>Note: The <code>dsenv</code> file does not work on RHEL7 and SLES12 systems, because the <code>\$DSHOME</code> variable is not exported correctly. Instead, you can use the <code>dsenv.multi.u</code> file.</p> <p>For example:</p> <pre data-bbox="602 1034 1220 1055">/opt/IBM/InformationServer/Server/DSEngine/dsenv.multi.u</pre>
UseSystemD	<p>SystemD is a system and a service manager for Linux operating systems. It helps manage applications across Linux distributions that support the SystemD feature.</p> <p>When the UseSystemD attribute of the InfoSphere agent is enabled on SystemD-compliant systems, the agent starts the InfoSphere resource in the <code>system.slice</code>. When this attribute is disabled, a typical online function starts the resource in the <code>user.slice</code>.</p> <p>Default value: 0</p> <p>Example: 1</p>

Table 3-2 lists the optional attributes for the InfoSphere agent.

Table 3-2 Optional attributes

Attribute	Description
HostName	<p>The virtual host name for this InfoSphere instance. The client side components use this host name to connect to the server side components.</p> <p>This attribute is used for in-depth monitoring of InfoSphere DataStage components.</p> <p>Default Value: ""</p> <p>Example: localhost, myhostname.mydomain.com, 92.168.1.100</p>
Port	<p>The port number that is allocated to a specific engine tier component.</p> <p>This attribute is used for in-depth monitoring of InfoSphere DataStage components.</p> <p>Default Value: ""</p> <p>Example: 31538 (for DSEngine Port)</p>
MonitorProgram	<p>Absolute path name of an external, user-supplied monitor executable.</p> <p>For information about setting this attribute: See "Executing a customized monitoring program" on page 24.</p> <p>Default Value: ""</p> <p>Example 1.: ServerRoot/bin/myMonitor.pl</p> <p>Example 2.: ServerRoot/bin/myMonitor.sh arg1 arg2</p>
LevelTwoMonitorFreq	<p>This type-level attribute specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring. You can also override the value of this attribute at the resource level. The value indicates the number of monitor cycles after which the agent will monitor the InfoSphere DataStage component in detail.</p> <p>For example, the value 5 indicates that the agent will monitor the InfoSphere DataStage component in detail after every five online monitor intervals.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p>

Table 3-2 Optional attributes (*continued*)

Attribute	Description
IMF	This type-level attribute determines if the agent must perform intelligent resource monitoring. You can also override the value of this attribute at the resource level. See “About the keys of the IMF attribute” on page 22.
IMFRegList	An ordered list of attributes whose values are registered with the IMF notification module. The attribute values can be overridden at the resource level.

About the keys of the IMF attribute

The IMF type-level attribute uses the following keys:

Table 3-3 IMF attribute keys

Key	Description
Mode	<p>Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows:</p> <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources. <p>Note: The agent for IBM InfoSphere DataStage supports intelligent resource monitoring for online resources only. Hence, Mode should be set to either 0 or 2.</p> <p>Default: 2</p>

Table 3-3 IMF attribute keys (*continued*)

Key	Description
MonitorFreq	<p>This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring.</p> <p>If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources <p>Default: 5</p>
RegisterRetryLimit	<p>If you enable intelligent resource monitoring, the agent invokes the <code>imf_register</code> agent function to register the resource with the AMF kernel driver.</p> <p>The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes.</p> <p>Default: 3</p>

Attributes used in different resource configurations

For each resource configuration, some attributes may be used by the agent and others may not be used. Use the following tables to figure out which attributes must be configured for your resource depending on the required configuration for your resource.

In these tables, the following conventions hold true:

- "Yes" implies that the attribute is mandatory for the given configuration.
- "Opt" implies that configuring the attribute is optional for the given configuration.
- "-" implies that the attribute is not used by the agent for the given configuration.

[Table 3-4](#) shows the attributes used in different resource configurations.

Table 3-4 Attributes used in different resource configurations

Attribute	DSEngine	ASBNodeAgent	ASBLoggingAgent
ResLogLevel	Yes	Yes	Yes
Component	Yes	Yes	Yes
User	Yes	Yes	Yes
InfoSphereInstanceDir	Yes	Yes	Yes
InfoSphereInstance	Yes	-	-
EnvFile	Yes	-	-
Hostname	Opt	Opt	Opt
Port	Opt	Opt	Opt
MonitorProgram	Opt	Opt	Opt
LevelTwoMonitorFreq	Opt	Opt	Opt

Executing a customized monitoring program

You can configure the monitor function to execute a custom monitor utility to perform a user-defined InfoSphere server state check. The utility is executed in the context of the UNIX user that is defined in the User attribute. The environment is set by sourcing the file specified in the EnvFile attribute.

The monitor function executes the utility specified in the MonitorProgram attribute if the following conditions are satisfied:

- The specified utility is a valid executable file.
- The first-level process check indicates that the InfoSphere instance is online.
- The LevelTwoMonitorFreq attribute is either set to 0 or 1, and the second-level check indicates that the InfoSphere instance is online.
- The LevelTwoMonitorFreq attribute is set to greater than 1, but the second-level check is deferred for this monitoring cycle.

The monitor function interprets the utility exit code as follows:

110 or 0	InfoSphere server instance is online
100 or 1	InfoSphere server instance is offline

99	InfoSphere server instance is unknown
Any other value	InfoSphere server instance is unknown

To ensure that the custom monitor utility is always available to the agent, Veritas recommends storing the file in the directory where the InfoSphere server is installed.

Enabling and disabling intelligent resource monitoring manually

The intelligent resource monitoring feature is enabled by default. Review the following procedures to enable or disable intelligent resource monitoring manually.

The IMF resource type attribute determines whether an IMF-aware agent must perform intelligent resource monitoring.

To enable intelligent resource monitoring manually

- 1 Make the VCS configuration writable.

```
# haconf -makerw
```
- 2 Run the following command to enable intelligent resource monitoring of online resources:

```
# hatype -modify InfoSphere IMF -update Mode 2
```
- 3 If required, change the values of the MonitorFreq key and the RegisterRetryLimit key of the IMF attribute.
- 4 Save the VCS configuration.

```
# haconf -dump -makero
```
- 5 Restart the agent. Run the following commands on each node.

```
# haagent -stop agent_name -force -sys sys_name  
# haagent -start agent_name -sys sys_name
```

To disable intelligent resource monitoring manually

- 1 Make the VCS configuration writable.

```
# haconf -makerw
```
- 2 To disable intelligent resource monitoring for all the resources of a certain type, run the following command:

```
# hatype -modify InfoSphere IMF -update Mode 0
```

- 3** To disable intelligent resource monitoring for a specific resource, run the following command:

```
# hares -override resource_name IMF  
  
# hares -modify resource_name IMF -update Mode 0
```

- 4** Save the VCS configuration.

```
# haconf -dump -makero
```

Troubleshooting the agent for IBM InfoSphere DataStage

This chapter includes the following topics:

- [Using the correct software and operating system versions](#)
- [Meeting prerequisites](#)
- [Reviewing error log files](#)
- [Troubleshooting the configuration for IMF](#)

Using the correct software and operating system versions

Ensure that you use correct software and operating system versions.

For information on the software versions that the agent for IBM InfoSphere DataStage supports, see the Veritas Services and Operations Readiness Tools (SORT) site: <https://sort.veritas.com/agents>.

Meeting prerequisites

Before installing the agent for IBM InfoSphere DataStage, ensure that the following prerequisites are met.

For example, you must install the ACC library on VCS before installing the agent for IBM InfoSphere DataStage.

See [“Before you install the Cluster Server agent for IBM InfoSphere DataStage”](#) on page 12.

Reviewing error log files

If you face problems while using InfoSphere or the agent for IBM InfoSphere DataStage, use the log files described in this section to investigate the problems.

Using InfoSphere log files

If an instance of the InfoSphere server component faces problems, you can access the server log files to further diagnose the problem. The log files are located in the following directories:

- `<ISInstallDir>/Server/PXEngine/java`
- `<ISInstallDir>/logs`
- `<ISInstallDir>/ASBNode/logs`
- `<ISInstallDir>/ASBNode/install/logs`

Reviewing cluster log files

In case of problems while using the agent for IBM InfoSphere DataStage, you can access the engine log file for more information about a particular resource. The engine log file is located at `/var/VRTSvcs/log/engine_A.log`.

For a long running cluster, the log files are rotated as `engine_B.log`, `engine_C.log`, and so on. The most-recent engine logs are present in the `engine_A.log` file.

Using trace level logging

The `ResLogLevel` attribute controls the level of logging that is written in a cluster log file for each InfoSphere resource. You can set this attribute to `TRACE`, which enables very detailed and verbose logging.

If you set `ResLogLevel` to `TRACE`, a very high volume of messages are produced. Veritas recommends that you localize the `ResLogLevel` attribute for a particular resource.

The `LogDbg` attribute should be used to enable the debug logs for the ACCLib-based agents when the ACCLIB version is 6.2.0.0 or later and the VCS version is 6.2 or later.

To localize ResLogLevel attribute for a resource

- 1 Identify the resource for which you want to enable detailed logging.
- 2 Localize the ResLogLevel attribute for the identified resource:

```
# hares -local Resource_Name ResLogLevel
```

- 3 Set the ResLogLevel attribute to TRACE for the identified resource:

```
# hares -modify Resource_Name ResLogLevel TRACE -sys SysA
```

- 4 Test the identified resource. The function reproduces the problem that you are attempting to diagnose.

- 5 Set the ResLogLevel attribute back to INFO for the identified resource:

```
# hares -modify Resource_Name ResLogLevel INFO -sys SysA
```

- 6 Save the configuration changes.

```
# haconf -dump
```

- 7 Review the contents of the log file.

You can also contact Veritas support for more help.

To enable debug logs for all resources of type InfoSphere

- ◆ Enable the debug log.

```
# hatype -modify InfoSphere LogDbg DBG_5
```

To override the LogDbg attribute at resource level

- ◆ Override the LogDbg attribute at the resource level and enable the debug logs for the specific resource.

```
# hares -override InfoSphere LogDbg
# hares -modify InfoSphere LogDbg DBG_5
```

Troubleshooting the configuration for IMF

If you face problems with the IMF configuration or functionality, consider the following:

- Ensure that the following attributes are configured with appropriate values.
 - AgentFile

- IMF
- IMFRegList
 - If IMFRegList is not configured correctly, the IBM InfoSphere DataStage resources that have been registered for IMF get unregistered every time the monitor function is run.
- If you have configured the required attributes to enable the IBM InfoSphere DataStage agent for IMF, but the agent is still not IMF-enabled, restart the agent. The `imf_init` function runs only when the agent starts up, so when you restart the agent, `imf_init` runs and initializes the IBM InfoSphere DataStage agent to interface with the AMF kernel driver.

- You can run the following command to check the value of the MonitorMethod attribute and to verify that a resource is registered for IMF.

```
# hares -value resource MonitorMethod system
```

The MonitorMethod attribute specifies the monitoring method that the agent uses to monitor the resource:

- Traditional—Poll-based resource monitoring
- IMF—Intelligent resource monitoring
- You can use the `amfstat` to see a list of registered PIDs for an InfoSphere resource.

The `amfstat` command shows the PIDs monitored by the InfoSphere DataStage agent.

```
[root@vcslx202 tmp]# amfstat
AMF Status Report
```

```
Registered Reapers (2):
```

```
=====
RID      PID      EVENT    REAPER
50       4085     3        0      InfoSphere
```

```
Process ONLINE Monitors (3):
```

```
=====
RID      R_RID    PID      GROUP
51       50       21453    infosphere_res
53       50       4516     infosphereASB_res
56       50       4363     infospherelog_res
```

The agent identifies the process for an InfoSphere DataStage component by applying pattern matching on the output of the `ps -ef` command. The patterns for the different InfoSphere DataStage component processes are:

```
DSEngine          root 21453 1 0 Aug09
                  00:00:00 /opt/IBM/InformationServer/Server/DSEngine
                  /bin/dsrpcd

ASBNodeAgent      root 4516 4510 1 13:17
                  00:00:11 /opt/IBM/InformationServer/ASBNode/apps
                  /jre/bin/java...asb.agent..

ASBLoggingAgent   root 4363 1 0 13:16
                  00:00:02 /opt/IBM/InformationServer/ASBNode/apps
                  /jre/bin/java ...
                  logging.agent.LoggingAgentSocketImpl
```

- Run the following command to set the `ResLogLevel` attribute to `TRACE`. When you set `ResLogLevel` to `TRACE`, the agent logs messages in the `InfoSphere_A.log` file.

```
# hares -modify ResourceName ResLogLevel TRACE
```

- Run the following command to view the content of the AMF in-memory trace buffer.

```
# amfconfig -p dbglog
```

Known issues

This release of the agent for IBM InfoSphere DataStage has the following known issues:

Problem

An error message might appear when you run the `hares -offline` command to take a resource offline.

Description

When a resource is taken offline, it is unregistered from the AMF module. However, the `imf_register` function attempts to unregister the resource again.

This results in an error message from the engine log.

```
imf_register entry point failed with exit code(1)
```

Workaround

It is safe to ignore this error message.

Sample Configurations

This appendix includes the following topics:

- [About sample configurations for the agents for IBM InfoSphere DataStage](#)
- [Sample agent type definition for InfoSphere](#)
- [Sample configuration files](#)
- [Sample service group configurations](#)

About sample configurations for the agents for IBM InfoSphere DataStage

The sample configuration graphically depicts the resource types, resources, and resource dependencies within the service group. Review these dependencies carefully before configuring the agents for IBM InfoSphere DataStage. For more information about these resource types, refer to the *Cluster Server Bundled Agents Reference Guide*.

Sample agent type definition for InfoSphere

The sample agent type definition for InfoSphere is as follows:

```
type InfoSphere (
    static boolean AEPTIMEOUT = 1
    static str AgentFile = "/opt/VRTSvcs/bin/Script51Agent"
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/InfoSphere"
    static str ArgList[] = { ResLogLevel, State, IState, Component,
        User, InfoSphereInstanceDir, EnvFile, InfoSphereInstance,
        Hostname, Port, MonitorProgram, UseSystemD }
    str ResLogLevel = INFO
```

```
str Component
str User
str InfoSphereInstanceDir
str EnvFile
str InfoSphereInstance
str Hostname
int Port
str MonitorProgram
boolean UseSystemD = 0
)
```

Sample configuration files

A sample main.cf file is as follows:

```
include "types.cf"
include "InfoSphereTypes.cf"

cluster infosphere_clus (
    UserNames = { admin = chiGheHe }
    Administrators = { admin, a }
)

system NodeA(
)

system NodeA (
)

group InfoSphere_grp (
    SystemList = { NodeA = 0, NodeB = 1 }
)

InfoSphere infosphereASB_res (
    ResLogLevel = TRACE
    Component = ASBNodeAgent
    User = root
    InfoSphereInstanceDir = "/opt/IBM/InformationServer"
    Hostname = localhost
    Port = 31532
)

InfoSphere infosphere_res (
```

```
ResLogLevel = TRACE
Component = DSEngine
User = dsadm
InfoSphereInstanceDir = "/opt/IBM/InformationServer"
EnvFile = "/opt/IBM/InformationServer/Server/DSEngine
/dsenv"
InfoSphereInstance = ade
Hostname = localhost
Port = 31539
)

InfoSphere infosphereelog_res (
    ResLogLevel = TRACE
    Component = ASBLoggingAgent
    User = root
    InfoSphereInstanceDir = "/opt/IBM/InformationServer"
    Hostname = localhost
    Port = 31533
)

Volume vol_res (
    Volume = infosphere_dg_vol
    DiskGroup = infosphere_dg
)

DiskGroup dg_res (
    DiskGroup = infosphere_dg_vol
)

Mount mount_res (
    MountPoint = "/opt/IBM/InformationServer"
    BlockDevice = "/dev/vx/dsk/infosphere_dg
/infosphere_dg_vol"
    FSType = ext4
    MountOpt = rw
    FsckOpt = "-y"
)

IP ip_res (
    Device = eth0
    Address = "10.209.76.140"
    NetMask = "255.255.252.0"
)

NIC nic_res (
    Device = eth0
```

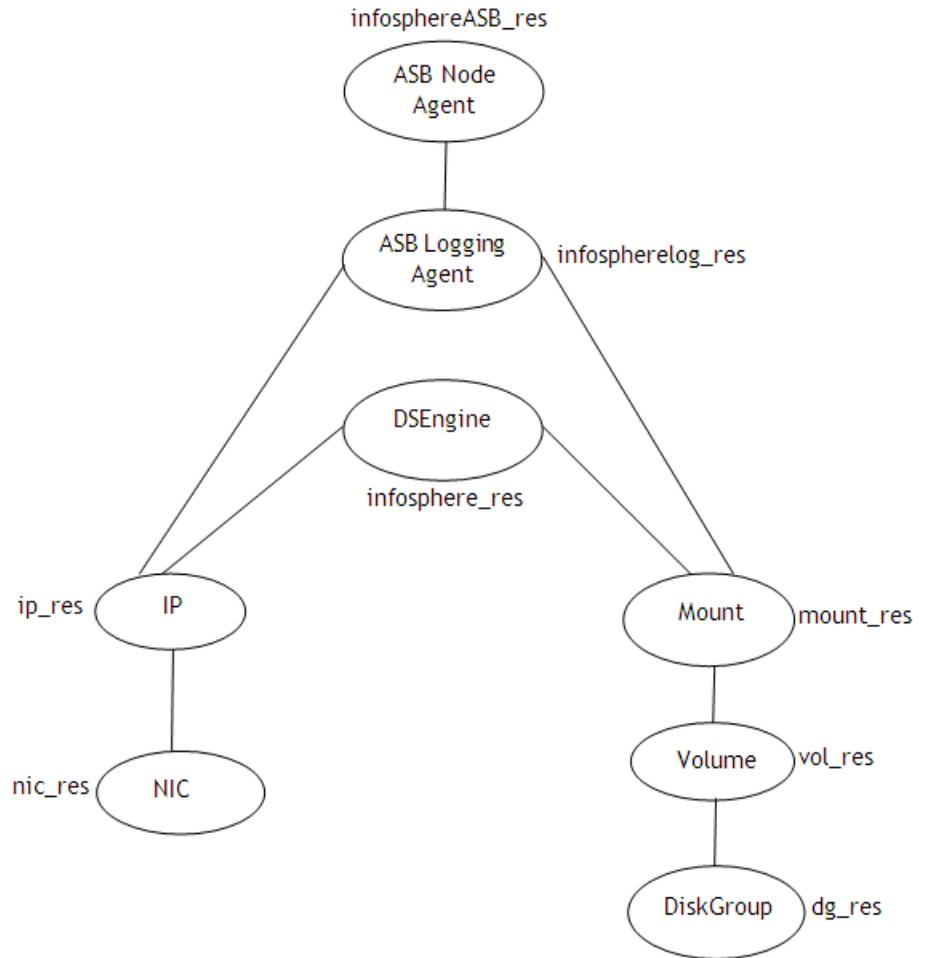
)

```
infosphereASB_res requires infosphereelog_res
infosphereelog_res requires mount_res
infosphere_res requires mount_res
infosphere_res requires ip_res
infosphereelog_res requires ip_res
ip_res requires nic_res
mount_res requires vol_res
vol_res requires dg_res
```

Sample service group configurations

Figure A-1 shows a service group with the ASB Node Agent, ASB Logging Agent, and DS Engine InfoSphere DataStage components running in a VCS environment.

Figure A-1 Sample service group for an InfoSphere instance



Sample resource configurations for InfoSphere

The following tables show the sample resource configurations for different InfoSphere DataStage components:

Table A-1 Sample resource configurations for DS Engine

Attribute	Sample Value
Component	DSEngine

Table A-1 Sample resource configurations for DS Engine *(continued)*

Attribute	Sample Value
User	Dsadm
Hostname	localhost
Port	31539
InfoSphereInstanceDir	/opt/IBM/InformationServer
InfoSphereInstance	ade
EnvFile	/opt/IBM/InformationServer/Server/DSEngine/dsenv

Table A-2 Sample resource configurations for ASB Node Agent

Attribute	Sample Value
Component	ASBNodeAgent
User	Root
Hostname	localhost
Port	31532
InfoSphereInstanceDir	/opt/IBM/InformationServer
InfoSphereInstance	-
EnvFile	-

Table A-3 Sample resource configurations for ASB Logging Agent

Attribute	Sample Value
Component	ASBLoggingAgent
User	Root
Hostname	localhost
Port	31533
InfoSphereInstanceDir	/opt/IBM/InformationServer

Table A-3 Sample resource configurations for ASB Logging Agent
(continued)

Attribute	Sample Value
InfoSphereInstance	-
EnvFile	-