



Cisco Tidal Enterprise Scheduler Hive Adapter Guide

Version: 6.2.1

May 4, 2016

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Tidal Enterprise Scheduler Hive Adapter Guide
© 2016 Cisco Systems, Inc. All rights reserved.



Preface	3
Audience	3
Related Documentation	3
Obtaining Documentation and Submitting a Service Request	3
Document Change History	4
Introducing the Hive Adapter	1-5
Overview	1-5
Prerequisites	1-6
Software Requirements	1-6
Configuring the Hive Adapter	2-7
Overview	2-7
Configuring the Hive Adapter	2-7
Licensing an Adapter	2-8
Securing the Adapter	2-8
Defining Runtime Users	2-9
Authorizing Schedulers to Work With Hive Jobs	2-10
Defining Scheduler Users of Hive Jobs	2-10
Defining a Connection	2-11
Adding a Connection	2-11
Using the Hive Adapter	3-15
Overview	3-15
Defining Hive Jobs	3-15
Hive Job Definition	3-15
Defining Hive Events	3-19
Hive Event Definition	3-19
Define an Action for an Event	3-21
Monitoring Hive Job Activity	3-22
Controlling Adapter and Agent Jobs	3-25
Holding a Job	3-25
Aborting a Job	3-26

Rerunning a Job	3-26
Making One Time Changes to an Adapter or Agent Job Instance	3-26
Deleting a Job Instance before It Has Run	3-26

Configuring service.props 4-29

About Configuring service.props	4-29
service.props Properties	4-29



Preface

This guide describes the installation, configuration, and usage of the Hive Adapter with Cisco Tidal Enterprise Scheduler (TES).

Audience

This guide is for administrators who install and configure the Hive Adapter for use with TES, and who troubleshoot TES installation and requirements issues.

Related Documentation

See the *Cisco Tidal Enterprise Scheduler Documentation Overview* for your release on cisco.com at:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/tidal-enterprise-scheduler/products-documentation-roadmaps-list.html>

...for a list of all TES guides.



Note

We sometimes update the documentation after original publication. Therefore, you should also review the documentation on Cisco.com for any updates.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>.

Subscribe to What's New in Cisco Product Documentation, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

Document Change History

The table below provides the revision history for the *Cisco Tidal Enterprise Scheduler Hive Adapter Guide*.

Version Number	Issue Date	Reason for Change
6.1.0	October 2012	<ul style="list-style-type: none">• New Cisco version.
6.2.1	June 2014	<ul style="list-style-type: none">• Available in online Help only.
6.2.1 SP2	June 2015	<ul style="list-style-type: none">• Configuration provided in the <i>TES Installation Guide</i>; usage provided in online Help only.
6.2.1 SP3	May 2016	<ul style="list-style-type: none">• Consolidated all Hive Adapter documentation into one document.



Introducing the Hive Adapter

This chapter provides an overview of the Hive Adapter and its requirements:

- [Overview](#)
- [Prerequisites](#)

Overview

The Cisco Tidal Enterprise Scheduler Hive Adapter provides the automation of HiveQL commands as part of the cross-platform process organization between Tidal Enterprise Scheduler (TES) and the TES Hadoop Cluster. The Adapter is designed using the same user interface approach as other Enterprise Scheduler adapter jobs, seamlessly integrating Hadoop Hive data management into existing operation processes.

The Hive Adapter allows you to access and manage data stored in the Hadoop Distributed File System (HDFS™) using Hive's query language, HiveQL. HiveQL syntax is similar to SQL standard syntax.

The Hive Adapter, in conjunction with Enterprise Scheduler, can be used to define, launch, control, and monitor HiveQL commands submitted to Hive via JDBC on a scheduled basis. The Adapter integrates seamlessly in an enterprise scheduling environment.

The Hive adapter includes the following features:

- Connection management to monitor system status with a live connection to the Hive Server via JDBC
- Hive job and event management includes the following:
 - Scheduling and monitoring of HiveQL commands from a centralized work console with Enterprise Scheduler
 - Dynamic runtime overrides for parameters and values passed to the HiveQL command
 - Output-formatting options to control the results, including table, XML, and CSV
 - Defined dependencies and events with Enterprise Scheduler for scheduling control
 - Runtime MapReduce parameters overrides if the HiveQL command results in a MapReduce job.

Prerequisites

- Hive version must be 0.9.0 or above.
- Hive Server
- The Hive Server must be fully operational and accessible to the Hive Adapter.

Cisco Tidal Enterprise Scheduler Adapters require Java 7. (Refer to *Cisco Tidal Enterprise Scheduler Compatibility Guide* for further details).

Software Requirements

The minimum requirement for the Enterprise Scheduler master and client is 6.1.0.

Refer to your *Cisco Tidal Enterprise Scheduler Compatibility Guide* for a complete list of hardware and software requirements.



Configuring the Hive Adapter

Overview

The Hive Adapter software is installed as part of a standard installation of Enterprise Scheduler. However, before the Hive Adapter can be used, the following configuration procedures must be completed:

- [Configuring the Hive Adapter](#) – Add the necessary properties to the service.props file.
- [Licensing an Adapter](#) – Apply the license to the Hive Adapter. You cannot define a Hive connection until you have applied the Hive license.
- [Securing the Adapter](#) – Define Hive users that the Adapter can use to establish authenticated sessions with the Hive server and permit requests to be made on behalf of the authenticated account.
- [Defining a Connection](#) – Define a Hive connection so the master can communicate with the Hive server.

See [Configuring service.props](#) for details about configuring service.props to control such things as polling, output, and log gathering.

Configuring the Hive Adapter

You must add necessary properties to the service.props file. There are optional parameters that can also be set as described here.

To configure the Hive adapter:

-
- | | |
|---------------|--|
| Step 1 | Stop the Master. |
| Step 2 | In the {207463B0-179B-41A7-AD82-725A0497BF42} directory, create a Config subdirectory. |
| Step 3 | If necessary, create the service.props file in the Config directory (see Configuring service.props). |
| Step 4 | In the service.props file, add the jarlib properties as follows: <ul style="list-style-type: none">a. For Apache 1.1.2, add: <code>jarlib=apache1.1.2</code>b. For Cloudera 4, add: <code>jarlib=cdh4</code>c. For MapR add: <code>jarlib=apache1.1.2</code> |
| Step 5 | (Optional) Modify the other properties in service.props as desired to control polling, output, and log gathering. See Configuring service.props . |

Step 6 Restart the Master.

Licensing an Adapter

Each TES Adapter must be separately licensed. You cannot use an Adapter until you apply the license file. If you purchase the Adapter after the original installation of TES, you will receive a new license file authorizing the use of the Adapter.

You might have a Demo license which is good for 30 days, or you might have a Permanent license. The procedures to install these license files are described below.

To license an Adapter:

Step 1 Stop the master:

Windows:

- a. Click **Start** and select **Programs>TIDAL Software>Scheduler>Master>Service Control Manager**.
- b. Verify that the master is displayed in the **Service** list and click on the **Stop** button to stop the master.

UNIX:

Enter **tesm stop**

Step 2 Create the license file:

- For a Permanent license, rename your Permanent license file to *master.lic*.
- For a Demo license, create a file called *demo.lic*, then type the demo code into the *demo.lic* file.

Step 3 Place the file in the **C:\Program File\TIDAL\Scheduler\Master\config** directory.

Step 4 Restart the master:

Windows:

Click **Start** in the Service Control Manager.

UNIX:

Enter **tesm start**

The master will read and apply the license when it starts.

Step 5 To validate that the license was applied, select **Registered License** from **Activities** main menu.

Securing the Adapter

There are two types of users associated with the Hive Adapter; **Runtime Users** and **Schedulers**. Although all connections to the Hive Server are anonymous, Enterprise Scheduler's job model requires at least one Hive user be defined. You maintain definitions for both types of users from the **Users** pane.

• Runtime Users

Hive Server connections are anonymous, but Enterprise Scheduler's job model requires at least one Hive runtime user. Therefore when defining the Hive runtime user, the password can be of any value as it is being used at runtime.

- **Schedulers**

Schedulers are those users who will define and/or manage Hive jobs. There are three aspects of a user profile that grant and/or limit access to scheduling jobs that affect Hive:

- Security policy that grants or denies add, edit, delete and view capabilities for Hive jobs and events.
- Authorized runtime user list that grants or denies access to specific authentication accounts for use with Hive jobs.
- Authorized agent list that grants or denies access to specific Hive Adapter connections for use when defining Hive jobs.

Defining Runtime Users

To define a runtime user:

-
- Step 1** From the **Navigator** pane, expand the **Administration** node and select **Runtime Users** to display the defined users.
- Step 2** Right-click **Runtime Users** and select **Add Runtime User** from the context menu (*Insert mode*).
-or-
Click the **Add** button on the menu bar.
The **User Definition** dialog displays.
- Step 3** Enter the new user name in the **User Name** field.
- Step 4** For documentation, enter the **Full Name** or description associated with this user.
- Step 5** In the **Domain** field, select a Windows domain associated with the user account required for authentication, if necessary.
- Step 6** To define this user as a runtime user for Hive jobs, click **Add** on the **Passwords** tab.
The **Change Password** dialog displays.

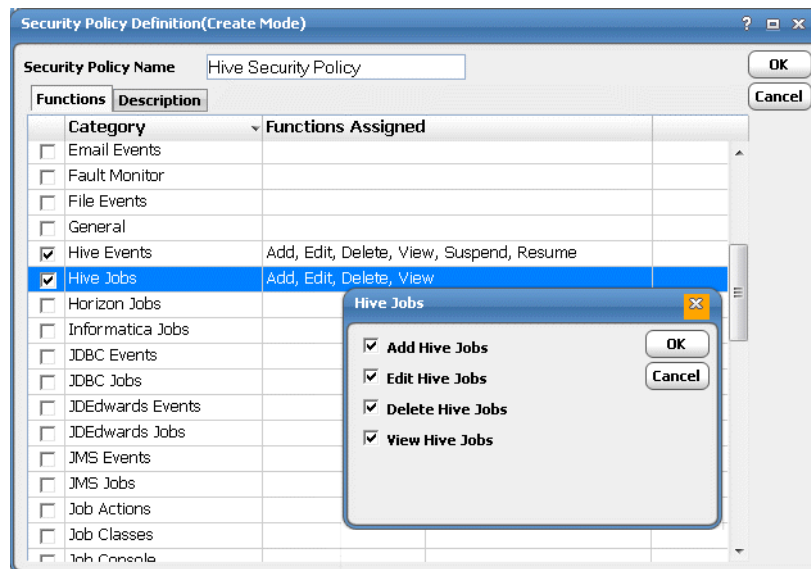
- Step 7** Select **Hive** from the **Password Type** list.
- Step 8** Enter a password (along with confirmation) in the **Password/Confirm Password** fields.
Since the password entered here is only used to satisfied the Enterprise Scheduler's job model but is not used to authenticate user to Hive Server. The password can be of any value in this case.
- Step 9** Click **OK** to return to the **User Definition** dialog.
The new password record displays on the **Passwords** tab.
- Step 10** Click **OK** to add or save the user record in the Enterprise Scheduler database.

For further information about the **User Definition** dialog, see your *Cisco Tidal Enterprise Scheduler User Guide*.

Authorizing Schedulers to Work With Hive Jobs

To authorize Schedulers:

- Step 1** From the **Navigator** pane, select **Administration>Security Policies** to display the **Security Policies** pane.
- Step 2** Right-click **Security Policies** and select **Add Security Policy** from the context menu. You can also right-click to select an existing security policy in the **Security Policies** pane and select **Edit Security Policy**.



- Step 3** In the **Security Policy Name** field, enter a name for the policy.
- Step 4** On the **Functions** page, scroll to the **Hive Jobs** category, click the ellipses on the right-hand side of the dialog and select the check boxes next to the functions that are to be authorized under this policy (**Add, Edit, Delete** and **View Hives Jobs**).
- Step 5** Click **Close** on the **Function** dropdown list.
- Step 6** Click **OK** to save the policy.

For further information about setting up security policies, see your *Cisco Tidal Enterprise Scheduler User Guide*.

Defining Scheduler Users of Hive Jobs

To define a Scheduler user to work with Hive jobs:

- Step 1** From the **Navigator** pane, expand the **Administrative** node and select **Interactive Users** to display the defined users.

- Step 2** Right-click **Interactive Users** and select **Add Interactive User** from the context menu (*Insert mode*). You can also right-click a user in the **Interactive Users** pane and select **Edit Interactive User** from the shortcut menu (*Edit mode*).
- The **User Definition** dialog displays.
- Step 3** If this is a new user definition, enter the new user name in the **User/Group Name** field.
- Step 4** For documentation, enter the **Full Name** or description associated with this user.
- Step 5** In the **Domain** field, select a Windows domain associated with the user account required for authentication, if necessary.
- Step 6** On the **Security** page, select the **Other** option and then select the security policy that includes authorization for Hive jobs.
- Step 7** Click the **Runtime Users** tab.
- Step 8** Select the Hive users that this scheduling user can use for Hive authentication from Hive jobs.
- Step 9** Click the **Agents** tab.
- Step 10** Select the check boxes for the Hive connections that this scheduling user can access when scheduling jobs.
- Step 11** Click the **Kerberos** page. If your Hadoop cluster is Kerberos secured, the Kerberos Principal and Kerberos Key page file is required. The Key page file is relative to the Master's file system and contains one or more Kerberos principals with their defined access to Hadoop.
- Step 12** Click **OK** to save the user definition.

Defining a Connection

You must create one or more Hive connections before Enterprise Scheduler can run your Hive jobs. These connections also must be licensed before Enterprise Scheduler can use them. A connection is created using the **Connection Definition** dialog.

Adding a Connection

To add a connection:

-
- Step 1** From the **Navigator** pane, navigate to **Administration>Connections** to display the **Connections** pane.
- Step 2** Right-click **Connections** and select **Add Connection>Hive Adapter** from the context menu. The **Hive Adapter Connection Definition** dialog displays.
- Step 3** On the **General** page, enter a name for the new connection in the **Name** field.
- Step 4** In the **Job Limit** field, select the maximum number of concurrent active processes that Enterprise Scheduler should submit to the Hive server at one time.
- Step 5** (Optional) From the **Default Runtime User** dropdown list, you can select the name of a default user for Hive jobs.

Only authorized users that have been defined with Hive passwords display in this list. The selected user is automatically supplied as the default runtime user in a new Enterprise Scheduler Hive job definition.

- Step 6** (Optional) Select the **Use as default for Hive Jobs** option if you wish for this connection to be used as the default whenever creating Hive jobs.
- Step 7** Click the **Hive Connection** tab.

- Step 8** In the **Hive Server JDBC URL** field, enter the URL of the Hive server in the format specified above.
- Step 9** Click the **Test** button to test connectivity using the specified Hive Server JDBC URL.
- If the test is unsuccessful, the following information dialog box displays:



- Step 10** Click the **Hive Parameters** tab to specify MapReduce job runtime parameters. These parameter values are applied if the HiveQL command results in a MapReduce job. Setting these parameter values at the connection level will globally set the job parameters for every job that is run on this connection. This configuration may be overridden at the job level by defining the corresponding configuration for the job options during the job definition.

Connection Definition (Create Mode) [[Hive]]

Hive Adapter

Name: Hive Connection

General | **Hive Connection** | Hive Parameters | Options | Description

Hive Connection Parameters

Name	Value
hive.archive.enabled	true
hive.auto.convert.join	true
hive.exec.max.dynamic.partitions	1000
hive.hashtable.initialCapacity	50000
hive.variable.substitute	false

Buttons: Add, Edit, Delete

☒ Enabled ☒ Use as default for Hive Jobs

Step 11 Click the **Options** tab, then in the **Event Poll** section, enter the number of minutes between successive poll checks of an event monitor that the Adapter uses to check whether it should raise an event. The default value is 60 minutes. Since MapReduce jobs are created and run each time a poll check occurs, it is recommended to keep the poll interval as high as possible.

Connection Definition (Edit Mode) [Hive-DEV-Extended[Hive]]

Hive Adapter

Name: Hive-DEV-Extended

General | **Hive Connection** | **Hive Parameters** | **Options** | Outages | Description

Event Poll

Poll Interval (in minutes): 60

Connection Attribute Queries

Attribute	SQL Query
Health	show databases
Local Time	select unix_timestamp() from default.hiveattributes limit 1
UTC Time	select from_unixtime(unix_timestamp()+3600*7) from default.hiveattribu
Version	select version from default.hiveattributes limit 1

Configuration Parameters

Name	Value
PrintStepHeader	Y

Buttons: Add, Edit, Delete

☐ Enabled ☐ Use as default for Hive Jobs

Step 12 In the **Connection Attribute Queries** section, define custom queries to provide additional information about the connection. Since these custom queries result in frequent running of MapReduce jobs at the Hadoop cluster in order to provide output, it is recommended these settings be skipped unless the information is vital to the running of jobs.

- **Version** – Queries the version number of the database.
- **Local Time** – Queries the time of the database's server.
- **UTC Time** – Queries the UTC time for the database server. The local time and the UTC are used to get the time zone and calculate the time difference between the database server and the scheduling master.

- **Health** – Queries the health of the database based on any custom SQL select statement.

If the query returns a result set, the connection is considered healthy (green). If an empty result set is returned, the connection is considered unavailable (red). If none is specified, the connection uses query “show databases” by default to check the health of the connection.

The following is an example Connection Attribute query:

- Create a custom table in Hive.

```
CREATE TABLE HIVEATTRIBUTES (VERSION STRING) ROW FORMAT DELIMITED;
```

- Prepare a datafile in `/tmp/version.txt` that includes the version.

```
0.9.0
```

- Load the data.

```
LOAD DATA LOCAL INPATH '/tmp/version.txt' OVERWRITE INTO TABLE HIVEATTRIBUTES;
```

- Local Time

```
select unix_timestamp() from HIVEATTRIBUTES limit 1;
```

- Version

```
select version from hiveattributes limit 1;
```

- UTC Time

```
select from_unixtime(unix_timestamp()) from hiveattributes limit 1;
```

Step 13 Click **OK** to save the new Hive connection.

The status light next to the connection indicates whether the Enterprise Scheduler Master is connected to the Hive server. If the light is green, the Hive server is connected.

A red light indicates that the master cannot connect to the Hive server. Hive jobs will not be submitted without a connection to the Hive server. You can only define jobs from the Client if the connection light is green.

If the light is red, you can test the connection to determine the problem. Right-click the connection and select **Test** from the shortcut menu. A message displays on the **Test Hive Connection** dialog describing the problem. Or, go to **Operator>Logs** to look for error messages associated with this connection.



Using the Hive Adapter

Overview

This chapter covers these topics:

- [Defining Hive Jobs](#)
- [Defining Hive Events](#)
- [Monitoring Hive Job Activity](#)
- [Controlling Adapter and Agent Jobs](#)

Defining Hive Jobs

This section provides instructions for defining a Hive job in Enterprise Scheduler and descriptions of the various types of tasks and options that can be included in the jobs.

Hive Job Definition

This section describes the basic steps for defining a Hive job.

To define a Hive job:

-
- Step 1** In the **Navigator** pane, select **Definitions>Jobs** to display the **Jobs** pane.
- Step 2** Right-click **Jobs** and select **Add>Hive Job** from the context menu.
- The **Hive Job Definition** dialog displays.

Hive Job Definition [Hive Job Example]

Hive Job Name:

Job Class:

Parent Group:

Owner:

Agent/Adapter Information

Agent/Adapter Name:

Agent List:

Runtime User:

Tracking

Use: ☒ Exit code
☐ External
☐ Scan output: Normal String(s)
☐ Scan output: Abnormal String(s)

Duration(in minutes)

Estimated: Minimum: Maximum:

☒ Exclude Completed Abnormally

☒ Enabled

Last Modified : 05/31/2012 10:50:31

The **Run** tab is selected by default. You must first specify a name for the job, the Hive Adapter connection that will be used for the job and a valid runtime user who has the appropriate authority for the batch job being scheduled.

Step 3 In the upper portion of the dialog, specify the following information to describe the job:

- **Job Name** – Enter a name that describes the job.
- **Job Class** – If you want to assign a defined job class to this job, select it from the dropdown list. This field is optional.
- **Owner** – Select the Hive job owner.
- **Parent Group** – If this job exists under a parent group, select the name of the parent group from the dropdown list. All properties in the Agent Information section are inherited from its parent job group.

Step 4 Specify the following connection information in the **Agent/Adapter Information** section:

- **Agent/Adapter Name** – Select the Hive Adapter connection to be used for this job from the dropdown list.
- **Agent List Name** – Select a list for broadcasting the job to multiple Hive servers rather than a specific agent/adapter.
- **Runtime User** – Select a valid runtime user with the appropriate authority for the job from the dropdown list.

Step 5 Specify the appropriate **Tracking** and **Duration** information for the job.

Refer to the *Cisco TES User Guide* for information on these options.

Step 6 Click the **Hive Job** tab.

Step 7 On the **HiveQL** tab, select a database from the **Database** drop-down list. By default, all commands are run on this “default” database.

Step 8 On the **HiveQL Commands** subtab, define one or more HiveQL commands.

If more than one command is defined (as in the figure above), separate them by a semicolon “;”. If parameters are necessary, precede them with a colon “:”.

If more than one databases specification is needed by the job definition, specify the additional databases using the <database> command to switch.

For example:

```
select * from sales_default;
use db01;
select * from sales_db01;
```

In the example above, the database is “Default” and the first statement is run against the “Default” database. The second command switches the current database to “db01”. All commands executed from this point forward are against “db01”.

Step 9 From the **Output Format** list, select one of the following query results formats.

- **XML** – Writes the query results in XML format.
- **Align Columns** – Displays the values in the most readable format.
- **CSV Format** – Separates values with commas.
- **Raw** – Separates values with a user-defined character.

Step 10 In the **Delimiter** field, specify the custom character to use for delimiting the column data from the query results if **Raw** is selected for **s**.

Step 11 (Optional) Select the **Include Headers** option to write out the column headers of the results.

You can view the steps on the **View Steps** subtab.

The screenshot shows the 'Hive Job Definition [Hive Job]' dialog box. At the top, there are input fields for 'Hive Job Name' (Hive Job), 'Job Class' (a dropdown), 'Parent Group' (a dropdown), and 'Owner' (qatest). Below these are tabs for 'Hive Job', 'Schedule', 'Run', 'Dependencies', 'Resources', 'Job Events', 'Options', 'Run Book', 'Notes', 'History', and 'Images'. The 'Hive Job' tab is active, showing a 'Database' dropdown and a 'HiveQL Commands' section with a table of commands. At the bottom, there are 'Output Format' (XML), 'Delimiter' (a text box), and 'Include Headers' (a checkbox). The 'Enabled' checkbox is checked, and the 'Last Modified' timestamp is 08/28/2012 19:51:33.

Step	HiveQL Command
1	select * from sales_default
2	select state, count(*) as count from sales_default
3	select state, sum(amount) as amount from sales_default

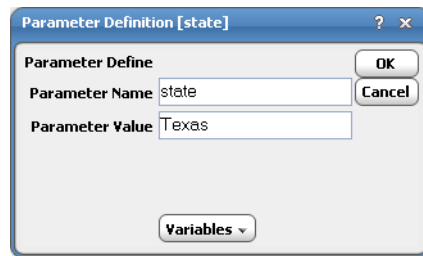
Step 12 Click the **Parameters** subtab to use parameters to be replaced at runtime.

The screenshot shows the 'Hive Job Definition [Hive Job Example]' dialog box with the 'Parameters' subtab selected. It shows a table with 'Name' and 'Value' columns. The 'state' parameter is highlighted with the value 'Texas'. An 'Edit' button is visible next to the table. The 'Enabled' checkbox is checked, and the 'Last Modified' timestamp is 05/31/2012 10:50:31.

Name	Value
state	Texas


This subtab provides a list of parameters that have been preceded by a colon in the HiveQL statements. If there are no parameters defined in the SQL statement, you cannot enter parameters on this subtab.

To edit a parameter, highlight it, then click **Edit** to display the **Parameters** dialog. Edit the parameter, then click **OK**.



You cannot modify the **Parameter Name**. The **Parameter Value** can be a literal value or a valid Enterprise Scheduler variable.

- Step 13** (Optional) On the **Options** subtab, specify additional MapReduce job runtime parameters. These parameter/values are applied if the HiveQL command results in a MapReduce job.



The parameter name/value specified at the job level overrides the value from the connection level if the same parameter name/value is specified.

- Step 14** Click **OK** to save the job.

Defining Hive Events

Using the Hive Adapter, you can define events that can be used for alerting and invoking an automated response through email and/or inserting additional jobs into the schedule. The **Event Definition** dialog is displayed when you add or edit a Hive event. Enterprise Scheduler can monitor events and then take one or more actions when the event trigger occurs. You must configure a calendar for the event from the **Schedule** tab to schedule when the event is enabled (that is, when monitoring will occur). If needed, you can configure the monitor to operate only during certain time periods or leave the monitor in operation at all times.

Hive Event Definition

To define a Hive event:

-
- Step 1** In the **Navigator** pane, select **Definitions>Events>Hive Events** to display the **Hive Events** pane.
- Step 2** Right-click **Hive Events** and select **Add Hive Event** from the context menus.
- The **Hive Event Definition** dialog displays.

Hive Event Definition (Create Mode)

Event Name: OK Cancel

Owner:

Hive Monitor | Schedule | Associated Action(s) | Description | Trigger History

Monitor

Hive Connection: ●

Hive Event:

Database Name:

Table Name:

Threshold:

☒ Public ☐ Enabled

Step 3 Enter a name for the event in the **Event Name** field and select an **Owner** from the dropdown list.

Step 4 Click the **Hive Monitor** tab to designate the data able in Hive to be monitored and define the condition that will be considered an event. Once the designated change to the datatable is detected, the event can be linked to an action to trigger an automatic response to the change.

Specify the following information:

- **Hive Connection** – Select the Hive connection from the list. This is the connection that will be monitored for the specified event. The status light next to the list indicates the connection status to the Hive server. If the light is green, the Hive server is connected.
- **Hive Event** – Select one of the following Hive events from the list:
 - **Row count exceed threshold** – When the threshold is exceeded and the event is triggered, the event triggers again if the row count drops below the threshold and then goes above the threshold again
 - **Column MAX value exceeds threshold** – If the largest monitored column value rises above the threshold, the event will trigger. The event will trigger again if the column maximum value drops below the threshold and rises above it again.
 - **Column MIN value below threshold** – If the smallest monitored column value falls below the threshold, the event will trigger. The event will trigger again if the column minimum value goes above the threshold and drops below it again.
 - **Custom query returns results** – A custom event triggers when a non-empty result set is returned from a custom query.
- **Database Name** – Select a database for the connection.
- **Table Name** – Select the name of the data table to be monitored.
- **Column Name** – This field displays if the event you have selected from the **Hive Event** list is either **Column MAX** or **Column MIN**. Select the name of the column for the selected table.
- **Threshold** – Select a threshold for the selected event.
- **Custom HiveQL** – This field displays if the event you have selected from the **Hive Event** list is **Custom query returns results**. Enter the HiveQL command.

**Note**

The other tabs on the Hive Event Definition dialog are general event configuration options and are not specific to the Hive Adapter. Any action that is available in Enterprise Scheduler, such as sending email, generating alerts, sending SNMP traps, setting variables, and adding jobs is available as a response to a Hive event.

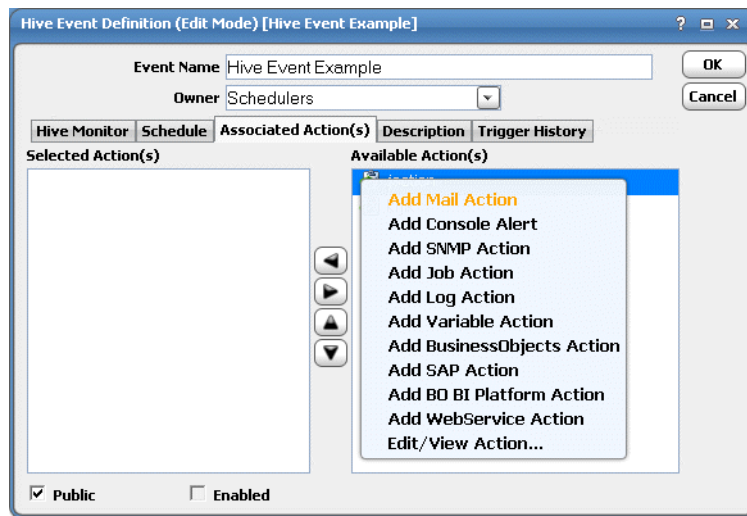
Step 5 Click **OK** to save the event definition.

Define an Action for an Event

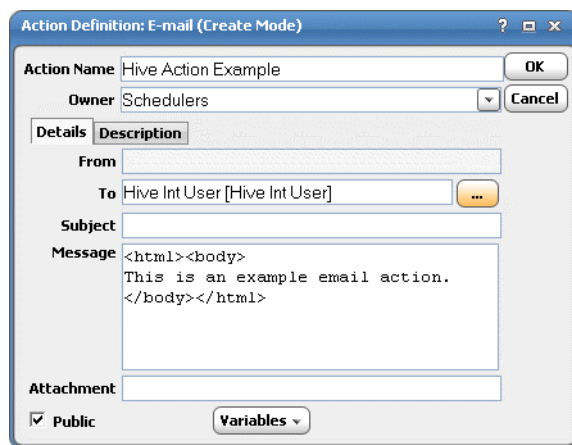
You can add any action for a Hive event that is available in Enterprise Scheduler.

For example, to define an email action for a Hive event:

-
- Step 1** From the **Hive Event Definition** dialog, click the **Associated Actions** tab.
- Step 2** In the **Available Actions** section, right-click and select **Add Mail Action**.



The **Action Definition: E-Mail** dialog displays.



Step 3 Complete the required fields on the **Action Definition** dialog and click **OK**.

Monitoring Hive Job Activity

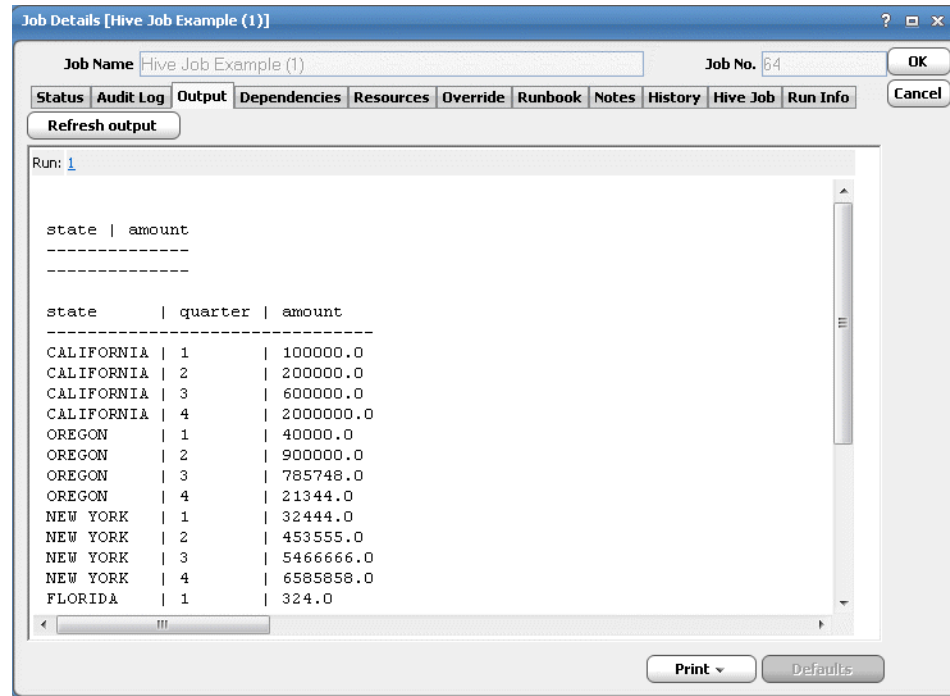
As Hive tasks run as pre-scheduled or event-based jobs, you can monitor the jobs as you would any other type of job in Enterprise Scheduler using the **Job Details** dialog. You can also use Business Views to monitor job activity and view when the jobs are active (see your *Cisco Tidal Enterprise Scheduler User Guide* for instructions on using Business Views).

To monitor job activity:

-
- Step 1** In the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.
 - Step 2** Right-click to select a job and choose **Details** from the context menu.

The **Job Details** dialog displays. On the **Status** page, you can view the status of the job, the start and end time, how long it ran, and how it was scheduled. The external ID is the run ID associated with the specific execution of the batch job.

- Step 3** Click the **Output** tab to view the result sets from the HiveQL commands after the job completes.

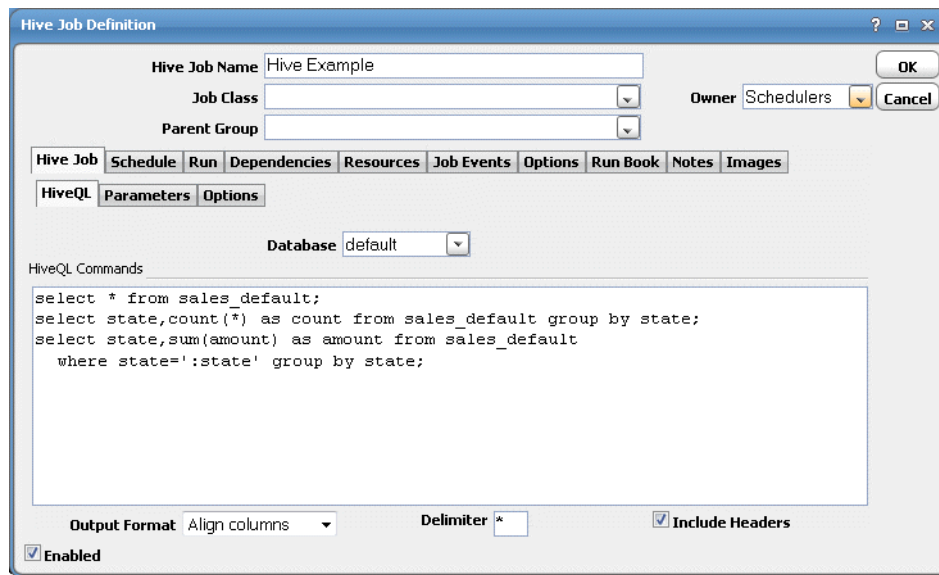


Enterprise Scheduler can be configured to save or discard job output by default from the **Defaults** tab of the **System Configuration** dialog. Regardless of the system default, any individual job instance can be configured from its job definition to override the system default. Each time a job is rerun, that run's output is separated by a block of number signs.

A Step Header with step number, status, and an abbreviated version of the step command is displayed before the actual result set. A Step Header can be turned off using the “PrintStepHeader” option in the connection. By default, PrintStepHeader is turned on.

If there are multiple commands in the Hive Job, each command is run as a separate step and all steps are displayed in the **Output** tab sequentially.

- Step 4** Click the **Hive Job** tab to view the original request along with variables used when this job was submitted. This tab allows you to override the HiveQL commands, output format, and parameter values prior to run or rerun. Overrides are not permitted when the job is running.



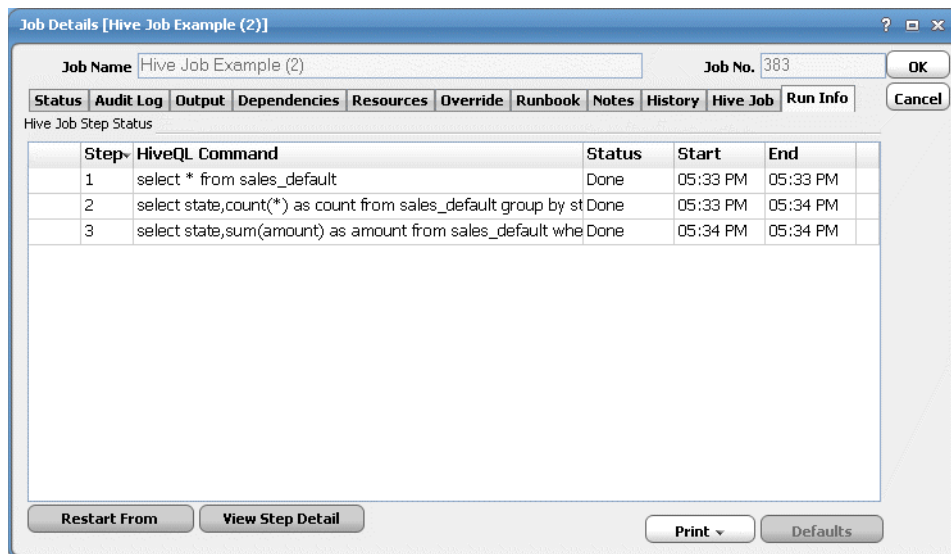
The **Hive Job Definition** dialog box is shown. It includes fields for **Hive Job Name** (Hive Example), **Job Class**, **Parent Group**, and **Owner** (Schedulers). There are tabs for **Hive Job** (Schedule, Run, Dependencies, Resources, Job Events, Options, Run Book, Notes, Images) and **HiveQL** (Parameters, Options). The **HiveQL** tab is active, showing a text area with the following commands:

```
select * from sales_default;
select state,count(*) as count from sales_default group by state;
select state,sum(amount) as amount from sales_default
  where state=':state' group by state;
```

Below the text area are settings for **Output Format** (Align columns), **Delimiter** (*), and **Include Headers** (checked). An **Enabled** checkbox is also present.

While the job is running, the fields are disabled; however, prior to running or rerunning the job, you can override any value on this screen. Your changes here only apply to this instance of the job (the original job definition is not affected).

- Step 5** Click the **Run Info** tab to view the run status, start, and end time for each step in the Hive Job. This tab is read-only.



The **Job Details [Hive Job Example (2)]** dialog box is shown. It includes fields for **Job Name** (Hive Job Example (2)) and **Job No.** (383). There are tabs for **Status**, **Audit Log**, **Output**, **Dependencies**, **Resources**, **Override**, **Runbook**, **Notes**, **History**, **Hive Job**, and **Run Info**. The **Run Info** tab is active, showing a table of Hive Job Step Status.

Step	HiveQL Command	Status	Start	End
1	select * from sales_default	Done	05:33 PM	05:33 PM
2	select state,count(*) as count from sales_default group by st	Done	05:33 PM	05:34 PM
3	select state,sum(amount) as amount from sales_default whe	Done	05:34 PM	05:34 PM

At the bottom are buttons for **Restart From**, **View Step Detail**, **Print**, and **Defaults**.

Click **View Details** to view the **Hive Job Run Info** dialog.

Hive Job Run Info

Close

Step: 2

Status: Done

Start: 12:00 AM

End: 12:00 AM

HiveQL Command:

```
select state, count(*) as count from sales_default group by state
```

Step Output:

state	count
CALIFORNIA	4
FLORIDA	4
NEW YORK	4
OHIO	4
OREGON	4

This dialog contains the run status, start, and end time for each step in the Hive job, the HiveQL command and its output. You can override the HiveQL command and output on this dialog.

Additionally, you can select a step and click Restart From to rerun a job from a particular step.

Step 6 When you have finished viewing the job activity details, click **OK** to close the dialog.

Controlling Adapter and Agent Jobs

Scheduler provides the following job control capabilities for either the process currently running or the job as a whole:

- [Holding a Job](#)—Hold a job waiting to run.
- [Aborting a Job](#)—Abort an active job.
- [Rerunning a Job](#)—Rerun a job that completed.
- [Making One Time Changes to an Adapter or Agent Job Instance](#)—Make last minute changes to a job.
- [Deleting a Job Instance before It Has Run](#)—Delete a job instance before it has run.

Holding a Job

Adapter/agent jobs are held in the same way as any other Scheduler jobs.

Adapter/agent jobs can only be held before they are launched. Once a job reaches the Adapter/Agent system, it cannot be held or suspended.

To hold a job:

-
- Step 1** From the **Job Activity** pane, right-click on the job.
- Step 2** Select **Job Control>Hold/Stop**.

Aborting a Job

Adapter/agent jobs are aborted in the same way as any other Scheduler jobs.

To abort a job:

-
- Step 1** From the **Job Activity** pane, right-click on the job.
- Step 2** Select **Job Control>Cancel/Abort**.

Rerunning a Job

On occasion, you may need to rerun an Adapter/Agent job. You can override parameter values first, if necessary, from the Adapter/Agent tab.

To rerun a job:

-
- Step 1** From the **Job Activity** pane, right-click the Adapter/Agent job you need to rerun.
- Step 2** Select **Job Control>Rerun** option from the context menu.

Making One Time Changes to an Adapter or Agent Job Instance

Prior to a run or rerun, you can edit data on the specific **Adapter/Agent** tab. To ensure that there is an opportunity to edit the job prior to its run, you can set the **Require operator release** option on the **Options** tab in the Adapter **Job Definition** dialog. Use this function to make changes to an Adapter job after it enters Waiting on Operator status as described in the following procedure.

To make last minute changes:

-
- Step 1** From the **Job Activity** pane, double-click the Adapter/Agent job to display the **Job Details** dialog.
- Step 2** Click the Adapter tab.
- Step 3** Make the desired changes to the job and click **OK** to close the **Job Details** dialog.
- Step 4** If this job is Waiting on Operator, perform one of the following tasks:
- To release the job, select **Job Control->Release**.
 - To rerun the job with changes, select **Job Control->Rerun**.

Deleting a Job Instance before It Has Run

Adapter/Agent job instances are deleted in the same way as any other Scheduler job.

Deleting a job from the **Job Activity** pane removes the job from the Scheduler job activity only. The original definition is left in tact.

To delete a job instance:

-
- Step 1** From the **Job Activity** pane, right-click the Adapter/Agent job to be deleted.
- Step 2** Select **Remove Job(s) From Schedule**.



Configuring service.props

About Configuring service.props

The **service.props** file is used to configure adapter behavior. **service.props** is located in the \config directory located under the Adapter's GUID directory. You can create both the directory and file if it does not yet exist. Properties that can be specified in service.props control things like logging and connection configuration. Many of the properties are specific to certain adapters; others are common across all adapters.

service.props Properties

The table below lists many of the parameters that can be specified in service.props. Some properties apply to all adapters (shaded in the table) and some properties are adapter-specific as indicated by the **Applicable Adapter(s)** column. The properties are listed in alphabetical order.

Property	Applicable Adapter(s)	Default	What It Controls
BYPASS_SEC_VALIDATION	Oracle Apps	N	If set to Y, the secondary user validation is bypassed. If not, secondary user validation is performed.
CLASSPATH	All	<none>	(Optional) – The path to the JDBC driver. If the default CLASSPATH used when the Adapter process is started does not include an appropriate JDBC driver jar required to connect to the PowerCenter Repository Database, you will need to specify this <i>service.props</i> configuration
CONN_SYNC	All	N	Setting this flag to Y allows synchronous connections without overloading the ROnly Thread. If set to N, the adapter might stop trying to reconnect after an outage or downtime.
DISCONN_ON_LOSTCONN	Informatica	N	Setting this flag to Y avoids an unnecessary logout call to the Informatica server when the connection is lost. This logout call usually hangs.

Property	Applicable Adapter(s)	Default	What It Controls
EnableDynamicPollingInterval	All	N	Use to avoid frequent polling on long-running jobs. When set to Y in service.props of a particular adapter, these properties are enabled: MinDynamicPollInterval—Minimum value should be 5 seconds. MaxDynamicPollIntervalInMin—Maximum value should be 5 minutes. PercentOfEstDuration—Default value is 5.
IGNORE_CODES	Informatica	<none>	This parameter can be set in service.props, job configuration and connection configuration parameters. The order of precedence is service.props (applicable for all jobs running in all connections), job level (only for that particular job), and connection (applicable for all jobs in the connection). This parameter is used to specify Informatica-specific error codes, separated by commas (,), that you want to ignore while running a job.
IGNORESUBREQ	Oracle Apps	N	Y or N. Setting this flag to Y stops huge job xml file transfers back and forth between the adapter and the AdapterHost during polls when a single request set has multiple sub-requests of more than 100. The default value is N or empty.
jarlib	Hive and MapReduce	<none>	Specifies the specific Java library to use for the adapter: <ul style="list-style-type: none"> For Apache 1.1.2, add: jarlib=apache1.1.2 For Cloudera 3, add: jarlib=cloudera For Cloudera 4, add: jarlib=cdh4 For MapR add: jarlib=apache1.1.2
kerbkdc	MapReduce	<none>	If the Hadoop cluster is Kerberos secured, use this value to specify the KDC Server. For example, kerbkdc=172.25.6.112
kerbrealm	MapReduce	<none>	If the Hadoop cluster is Kerberos secured, use this value to specify the Kerberos Realm. For example, kerbrealm=TIDALSOFT.LOCAL

Property	Applicable Adapter(s)	Default	What It Controls
Keystore	BusinessObjects , BusinessObjects BI, BusinessObjects DS, Cognos, JD Edwards, Oracle Applications, UCS Manager, VMware, Web Service	<none>	Specify Keystore=c:\\<adapter_certificate_directory>\\<your_trusted_keystore>.keystore when importing certificates into a Java keystore.
LAUNCH_DELAY (in milliseconds)	Informatica	<none>	This parameter can be set in service.props, job configuration and connection configuration parameters. The order of precedence is service.props (applicable for all jobs running in all connections), job level (only for that particular job), and connection (applicable for all jobs in the connection). If a non-zero value is set for this parameter, then the jobs are delayed for the specified number of milliseconds before being submitted to Informatica.
LoginConfig	BusinessObjects BI Platform, BusinessObjects Data Services	<none>	Specifies the location of the login configuration if using WinAD or LDAP authentication. For example: LoginConfig=c:\\windows\\bscLogin.conf where "c:\\windows\\bscLogin.conf" is the location of the login configuration information. Note the use of \\ if this is a Windows location.
MaxLogFiles	Informatica, JDBC	50	(Optional) – Number of logs to retain.
OUTPUT_ASYNC_LOGOUT	Informatica	N	Setting this flag to Y avoids jobs getting stuck in Gathering Output status.
OUTPUT_SYNC	All	Y	Enables concurrent output gathering on a connection. To enable this feature, set the value to N.
POLL_SYNC	All	Y	Enables concurrent polling on connections of the same type. This is helpful when there is a heavily load on one connection of an adapter. The heavily loaded connection will not affect the other adapter connection. To enable this feature, set the value to N.
QUERY_TIMEOUT	Oracle Apps	N	Y or N. If set to Y, the timeout value defined using the parameter QUERY_TIMEOUT_VALUE is applied to the SQL queries. Default value is N or empty.

Property	Applicable Adapter(s)	Default	What It Controls
QUERY_TIMEOUT_VALUE	Oracle Apps	unset	The time period in seconds that SQL queries wait before timeout. If 0 or not set, there is no timeout.
READPCHAINLOG	SAP	Y	Used to control the log gathering in SAP Process Chain jobs. This property depends on the Summary Only check box of the job definition Options tab.
SCANFOR_SESSIONSTATS	Informatica	Y	Y or N - Set this parameter to N to turn off the default behavior of Informatica jobs collecting the session statistics during the job run.
SCANFOR_SESSIONSTATS_AFTER_WF_ENDS	Informatica	N	Y or N - Set this parameter to Y to turn off the gathering of session statistics during each poll for the status of Informatica jobs.
TDLINFA_LOCALE	Informatica	<none>	Points to the Load Manager Library locale directory. See “Configuring the Informatica Adapter” in the <i>Informatica Adapter Guide</i> for how to set this for Windows and Unix environments.
TDLINFA_REQUESTTIMEOUT	Informatica	<none>	(Optional) – The number of seconds before an API request times out. The default is 120 seconds, if not specified.
TDLJDBC_LIBPATH	JDBC	<none>	(Windows only, optional) An alternate path to the JDBC library files. The library file path should have been configured given system environment variables. This option is available in case you wish to use an alternate set of libraries and may be helpful for trouble-shooting purposes.
TDLJDBC_LOCALE	JDBC	<none>	The path to the JDBC locale files.
TRANSACTION_LOG_BATCH_SIZE	MS SQL	5000	Set this parameter if more than 5000 lines need to be read from the transaction table.
version_pre898	JD Edwards	N	If running on a JD Edwards server version that is less than 8.9.8, set version_pre898=Y.