TIDAL

Tidal Workload Automation SSH Adapter Guide

Version 6.2.2

First Published: August 2018

tidalautomation.com

Contents

Preface

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

Audience

This guide is for administrators who install and configure the SSH Adapter for use with Cisco Tidal Enterprise Scheduler, and who troubleshoot TES installation and requirements issues.

Related Documentation

For a list of all Tidal Workload Automation guides, see the Tidal Workload Automation Documentation Overview of your release on tidalautomation.com at:

Tidal Automation Doc Support

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

Document Change History

The table below provides the revision history for the SSH Adapter Guide.

Version Number	Issue Date	Reason for Change
6.1.0	October 2012	• New Cisco version.
6.2.1	June 2014	• Available in online Help only.
6.2.1 SP2	June 2015	• Configuration provided in the *Cisco Tidal Enterprise Scheduler Installation Guide;* usage provided in online Help only.
6.2.1 SP3	May 2016	• Consolidated all SSH Adapter documentation into one document.

Introducing the SSH Adapter

This chapter provides an overview of the Cisco Tidal Enterprise Scheduler SSH Adapter and its requirements:

Overview

The Tidal Enterprise Scheduler Adapter for Secure Shell (SSH) is a network protocol that allows data to be exchanged using a secure channel between two networking devices. The SSH adapter allows Enterprise Scheduler to run commands or script activities on a system or network device that has SSH enabled.

In addition to the standard SSH scripts, the SSH adapter also allows you to execute commands against Cisco IOS network devices, such as the SSH Command, SSH Script, Get File and Put File activities.

This guide is intended to provide information on using the objects provided by the SSH Adapter. provides instructions for viewing SSH adapter properties, defining SSH targets and activities, instructions for completing the property pages for each specific activity, and instructions on viewing the activity results.

Prerequisites

Prior to configuring the SSH Adapter, you must ensure that the following prerequisites have been met.

Requirements

Enterprise Scheduler supports the following environment:

- Tidal Enterprise Scheduler 6.0 or above
- Version SSH-2
- Verify that SSHD is running on the SSH Server or another 3rd party SSH software.

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

Configuring the SSH Adapter

Overview

The SSH Adapter provides the ability to execute secure shell scripts and commands.

The SSH Adapter is installed as part of the Enterprise Scheduler. However, you must perform the following steps to license and configure the SSH adapter before you can run SSH jobs:

- Licensing an Adapter – License the SSH adapter. You cannot define a SSH connection until you have applied the SSH license from Tidal Software.
- Securing the SSH Adapter – Define a SSH Authentication user to authorize a connection to be established to the SSH agent and permit requests to be made on behalf of the authenticated account.
- Defining a SSH Adapter Connection – Define a SSH connection so the master can communicate with the SSH server.

See Configuring service.props for information about general and adapter-specific properties that can be set to control things like logging and connection properties.

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 SSH Adapter

Many operating system and application activities require credentials for proper execution. The Runtime Users feature is used to create a runtime user record to store the information about the user security context and to pass this information to the adapters for activity execution, event monitoring and some target operations (such as availability monitoring and discovery). When defining a process or certain activities, you can use the runtime user records that are defined in the product to assign credentials for the process or activity. There are two types of users associated with the SSH Adapter, Runtime Users and Schedulers. You maintain definitions for both types of users from the **Users** pane.

- **Runtime Users**

  Runtime users in the context of SSH jobs represent those users and passwords required for SSH Authentication. If the SSH server (the machine you are executing a job to) requires authentication based on user and password credentials, these users will need to be defined as runtime users.

- **Schedulers**

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

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

Defining Runtime Users

The credentials specified for the runtime user are used to store the information about the simple user security context consisting of a user name and password pair and to pass this information to the adapters. This runtime user can be used for database targets when needing database authentication.

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 Users** from the context menu (*Insert* mode). You can also right-click a user in the **Runtime Users** pane and select **Edit Runtime 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 To define this user as a runtime user for SSH jobs, click **Add** on the **Passwords** tab.

The **Change Password** dialog displays.

Step 7 Select **SSH** from the **Password Type** list.

Step 8 Enter a password (along with confirmation) in the **Password/Confirm Password** fields.

Only those users with a password specified for SSH will be available for use with SSH jobs. The password might be the same as the one specified for Windows/FTP jobs.

Step 9 Click **OK** to return to the **User Definition** dialog.

The new password record displays on the **Passwords** tab.

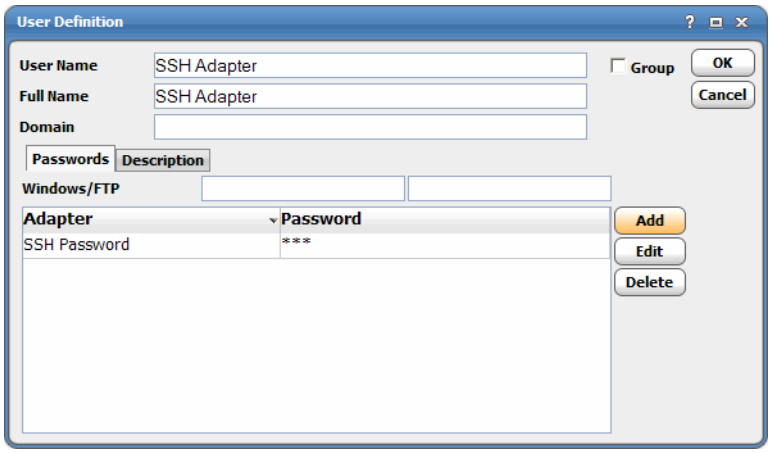

Step 10 Click **OK** to add or save the user record in the Enterprise Scheduler database.

Authorizing Schedulers to Work With SSH Jobs

Authorizing schedulers involves these tasks:

- Defining a Security Policy
- Defining SSH Scheduler Users

Defining a Security Policy

To define a Security Policy that authorizes access to SSH jobs:

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 a defined security policy in the **Security Policies** pane and select **Edit Security Policy**.

The **Security Policy Definition** dialog displays.

Note Refer to the *Tidal Enterprise Scheduler User Guide* for a general discussion on setting up security policies that you associate with Scheduler Users.

Step 3 In the **Security Policy Name** field, enter a name for the policy.

Step 4 On the **Functions** page, scroll to the **SSH Jobs** category, double-click the **Functions Assigned** field on the right-hand side of the dialog to view the **SSL Jobs** dialog.

Step 5 Select the check boxes next to the functions that are to be authorized under this policy (**Add**, **Edit**, **Delete** and **View SSH Jobs**).

Step 6 Click **OK** on the **SSL Jobs** dialog.

Step 7 Click **OK** to save the policy.

Defining SSH Scheduler Users

To define a Scheduler user to work with SSH jobs:

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

Step 2 Right-click **Interactive Users** and select **Add Users** 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.

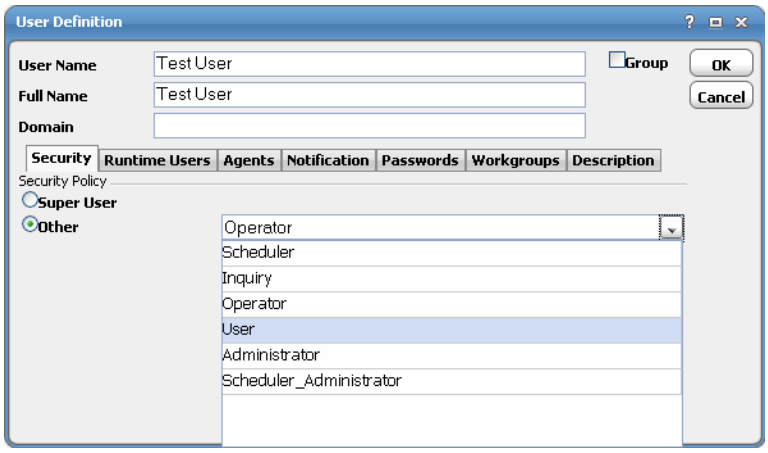

Step 3 If this is a new user definition, enter the new user name in the **User/Group Name** field.

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 SSH jobs.

Step 7 Click the **Runtime Users** tab.

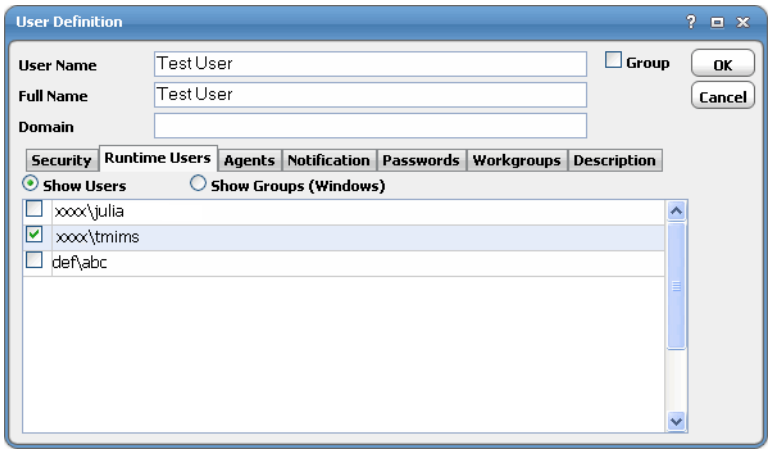

Step 8 Select the SSH users that this scheduling user may use for SSH authentication in SSH jobs.

Step 9 Click the **Agents** tab.

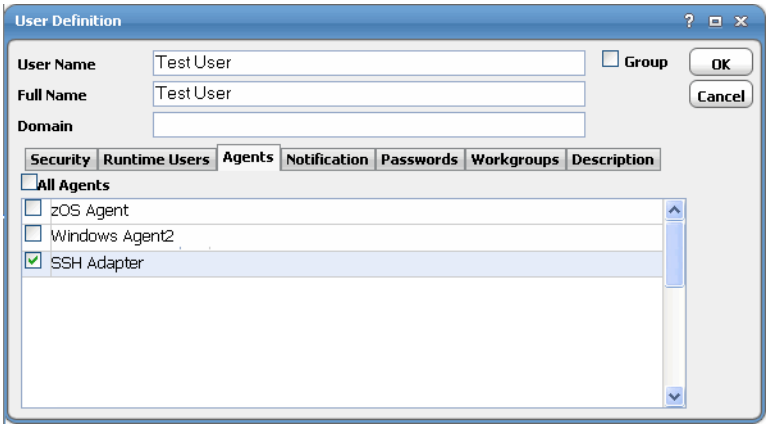

Step 10 Select the check boxes for the SSH connections that this scheduling user can access when scheduling jobs.

Step 11 Click **OK** to save the user definition.

Defining a SSH Adapter Connection

You must create a connection to a SSH server before Enterprise Scheduler can run your SSH jobs. These connections also must be licensed before Enterprise Scheduler can use them. A connection is created using the **Connection Definition** dialog.

Adding a SSH Adapter 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>SSH Adapter** from the context menu.

The **SSH Adapter Connection Definition** dialog displays.

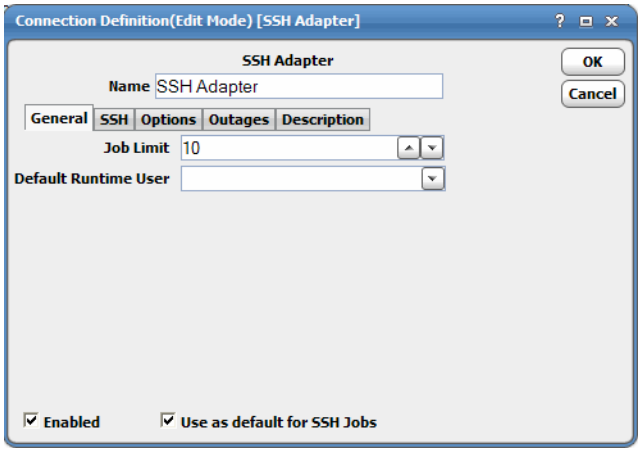

Step 3 On the **General** page, enter a name for the new connection in the **Name** field.

Step 4 In the **Job Limit** field, select the maximum number of concurrent active processes that Enterprise Scheduler should submit to the SSH server at one time.

Step 5 (Optional) From the **Default Runtime User** drop-down list, select the name of the default user for SSH jobs. The runtime user is used for SSH authentication and SSH uses this to authorize scheduled operations.

Only authorized users that have been defined with SSH passwords display in this list. The selected user is automatically supplied as the runtime user in Enterprise Scheduler SSH job definitions.

Step 6 Click the **SSH Connection** tab.

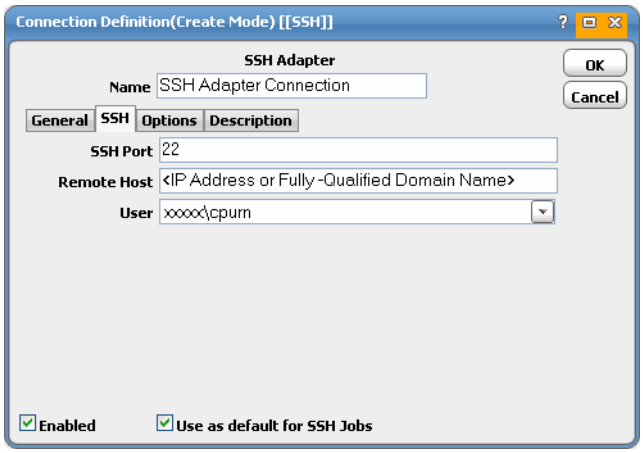

Step 7 In the **SSH Port** field, enter the appropriate port number for the SSH listener. The default port is `22`.

Note Port 22 must be open between Enterprise Scheduler and the SSH server.

Step 8 In the **Remote Host** field, enter the IP address or fully-qualified domain name for your target remote SSH host.

Step 9 (Optional) In the **User** field, select a user from the drop-down list who is authorized to connect and monitor attributes and invoke connection level operations.

Step 10 Click **OK**. The configured connection displays in the **Connections** pane.

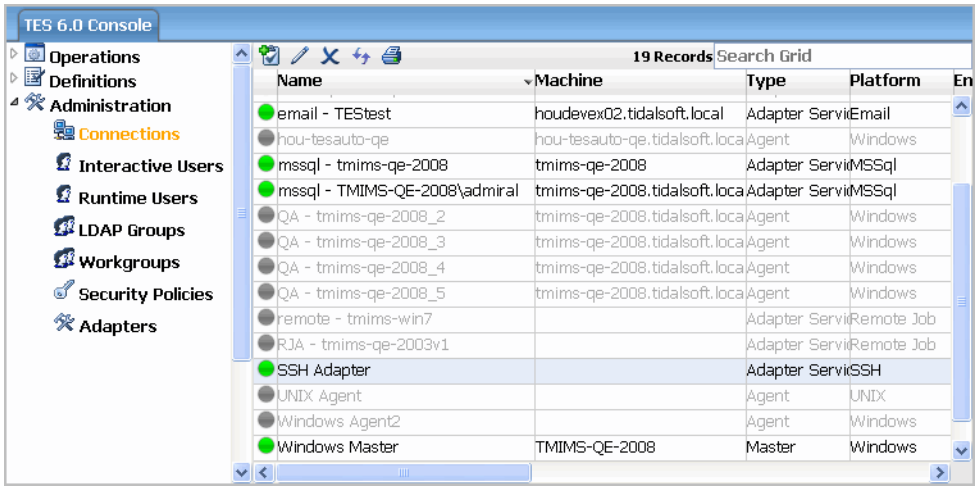

The status light next to the connection indicates whether the Enterprise Scheduler Master is connected to the SSH instance. If the light is green, the SSH instance is connected.

A red light indicates that the master cannot connect to the SSH instance. However, the jobs will not run without a connection to the SSH instance.

Note If there is an attribute associated with Health, this also determines whether the light is green or red.

If the light is red, check **Operations>Logs** for any associated error messages. You can also test the connection to determine the problem. Right-click the connection and select **Test** from the shortcut menu. A message displays on the **Test SSH Connection** dialog describing the problem.

Using the SSH Adapter

This chapter guides you through using the features of the SSH Adapter in Enterprise Scheduler, including:

- Defining SSH Jobs
- Monitoring SSH Jobs
- Controlling Adapter and Agent Jobs

Defining SSH Jobs

This section provides instructions for defining a SSH job in Enterprise Scheduler.

SSH Job Definition

To define a SSH job:

Step 1 In the **Navigator** pane, select **Definitions>Jobs** to display the **Jobs** pane.

Step 2 Right-click **Jobs** and select **Add>SSH Job** from the context menus.

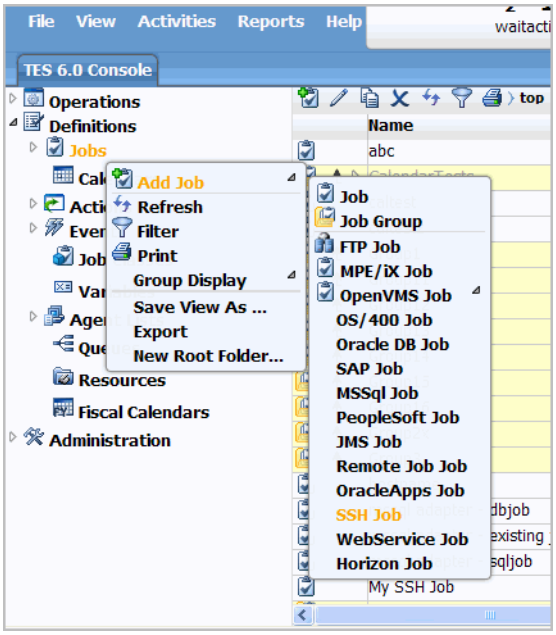

The **SSH Job Definition** dialog displays.

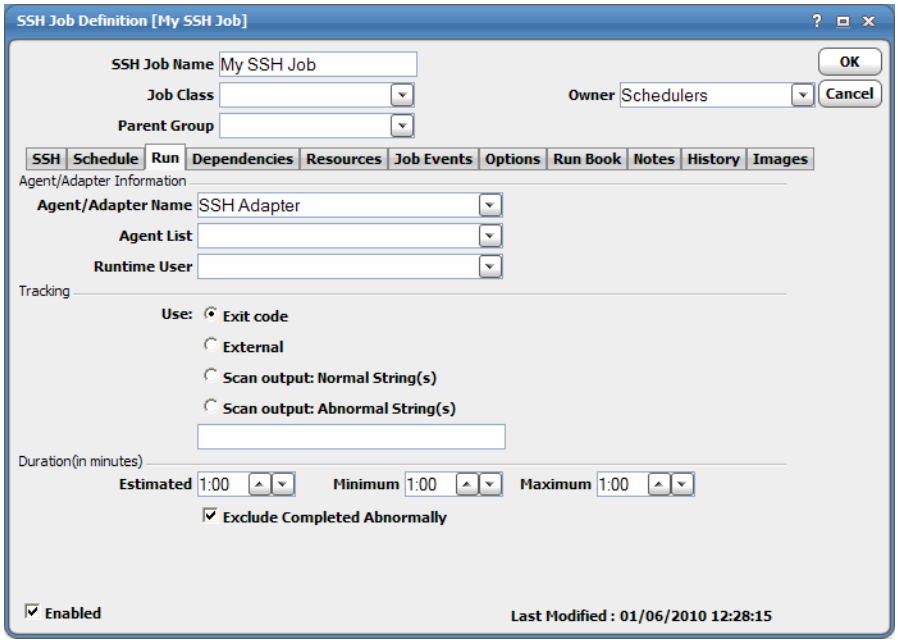

The **Run** tab is selected by default. You must first specify a name for the job, a valid runtime user who has the appropriate SSH authority for the operation, and the SSH adapter connection that will be used for the job.

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.

- (Optional) **Job Class** – If you want to assign a defined job class to this job, select it from the drop-down list.
- **Owner** – Select the user name from the drop-down list for the person who owns this job. The user must have the appropriate SSH authority for the operation.
- **Parent Group** – If this job exists under a parent group, select the name of the parent group from the drop-down 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 SSH adapter connection to be used for this job from the drop-down list.
- (Optional) **Runtime User** – Select a valid runtime user with the appropriate SSH authority for the job from the drop-down list.

Step 5 Specify the appropriate **Tracking** and **Duration** information for the job. Refer to the *Tidal Enterprise Scheduler User Guide* for information on these options.

Step 6 Click the **SSH** tab.

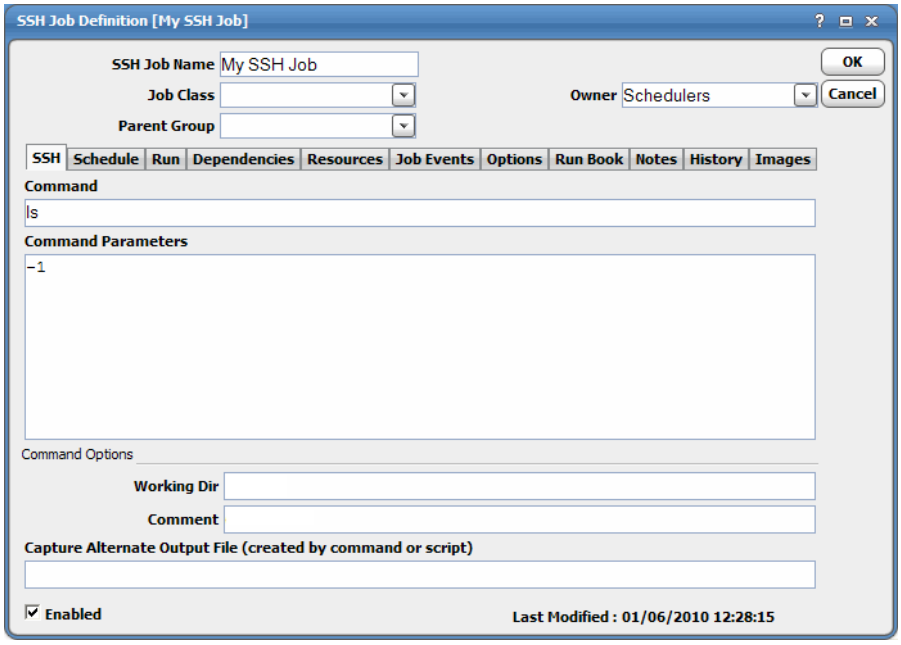

Step 7 In the **Command** field, enter the the absolute path and filename of the command, script, batch file or executable that you want the job to run.

Step 8 In the **Param** field, enter either the hard-coded value for each parameter or type a parameter name.

Step 9 In the **Working Directory** field, enter the path for the working directory of the program or script specified in the **Command** field.

Step 10 (Optional) In the **Comment** field, enter any comments.

Step 11 In the **Capture Alternate Output File** field, enter a filename.

If a file is specified, the SSH agent looks for the output file in the location specified and reads it as text, returning that text as the output of the job instead of gathering the job process's standard output. A job process's "standard output" is the text that would be returned to the screen if you were to run the process manually via a command line interface. However, some applications do not return any "standard output," resulting in a blank **Output** tab in the **Job Detail** dialog if run by Scheduler.

Step 12 Click **OK** to save the job.

Monitoring SSH Jobs

As SSH 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 the Business view to monitor job activity and view when the jobs are active (see the *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** console.

Step 2 Right-click job and select **Details** from the context menu.

The **Job Details** dialog displays.

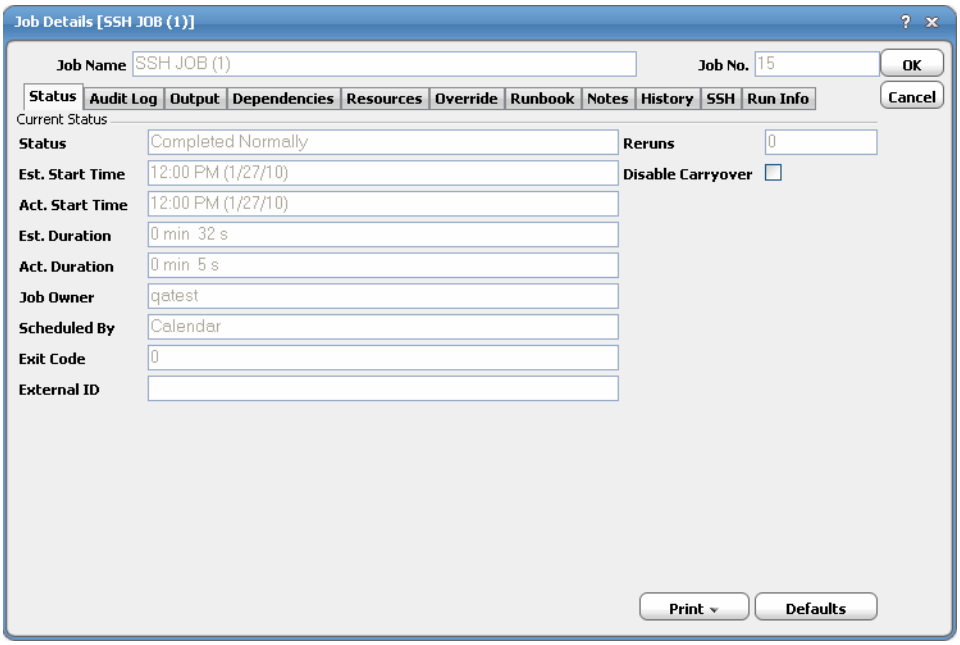

The **Status** page displays by default. You can view the status of the job, the start and end time, how long it ran, and how it was scheduled.

Step 3 Click the **Output** tab to view a task summary.

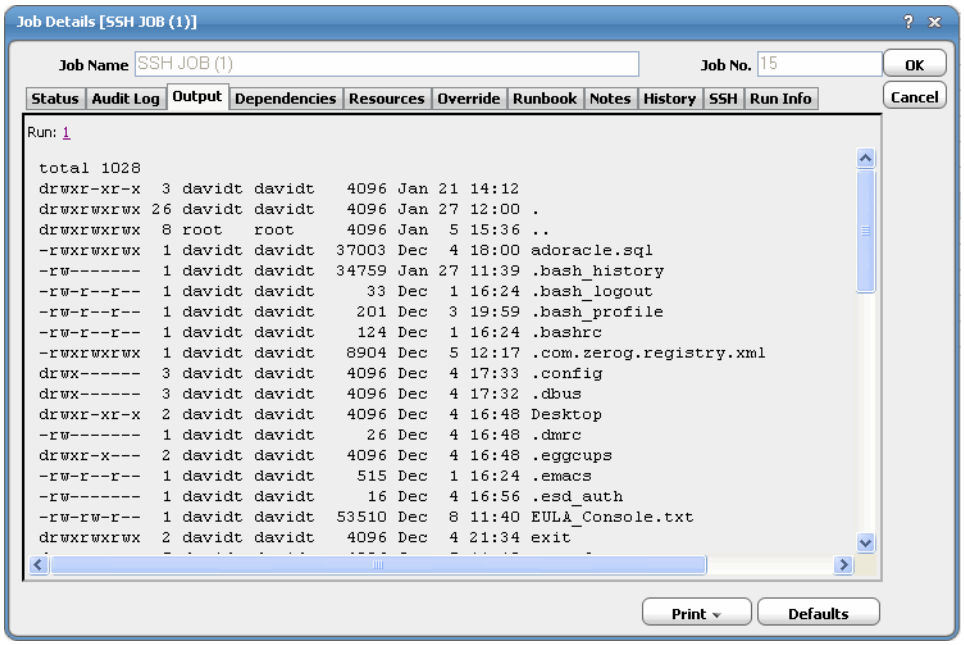

Step 4 Click the **Run Info** tab to view the request with the values used when this instance of the job was last run.

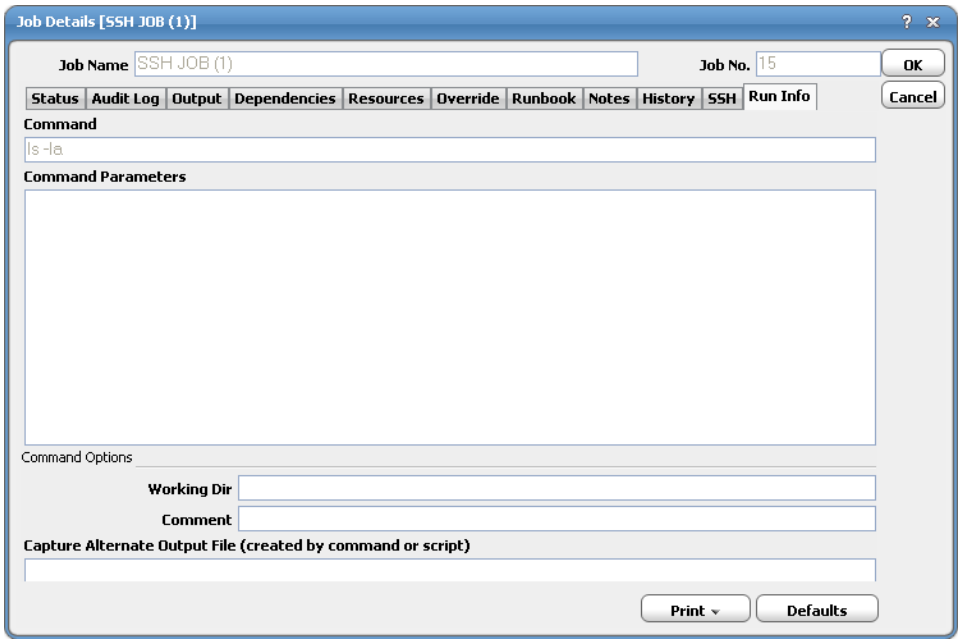

Step 5 Click the **SSH** tab to view the job definition details and the variables that were used when the job was submitted. Changes only affect this instance of the job and can only be made before the job runs the first time or prior to a rerun (not while the job is running).

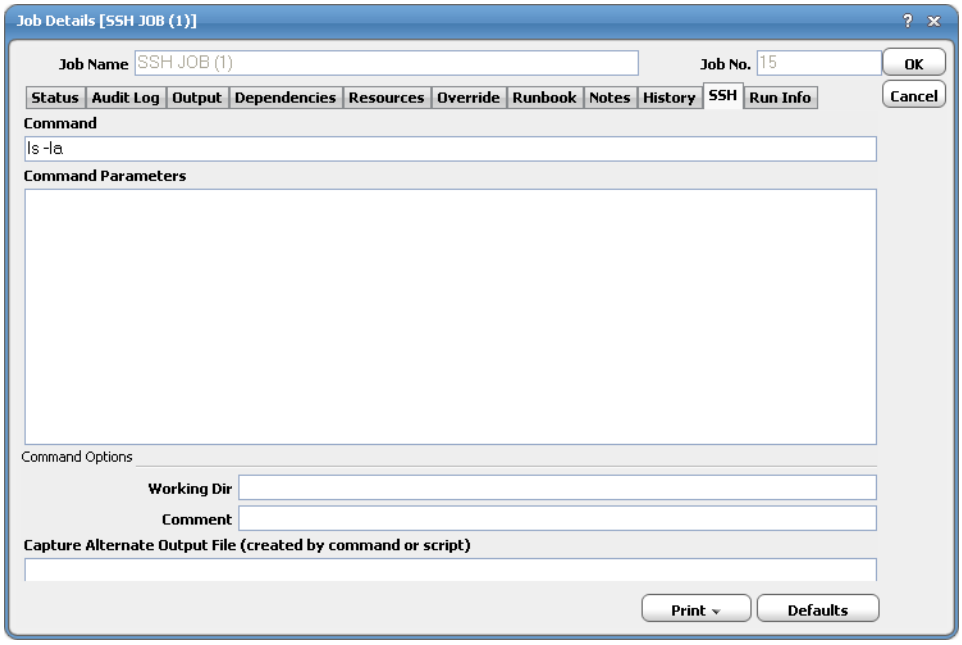

Step 6 When you have completed viewing the job activity details, click **OK** to close the dialog.

For jobs currently executing a custom job step, if there is a SSH operation that has been configured to suspend jobs, that operation is invoked. When the job is resumed, if there is a SSH operation that has been configured to resume suspended jobs, that operation is invoked.

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 *service.props* configuration
CONN_SYNC	All	N	Setting this flag to Y allows synchronous connections without overloading the RDOnly 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: o For Apache 1.1.2, add: **jarlib=apache1.1.2** o For Cloudera 3, add: **jarlib=cloudera** o For Cloudera 4, add: **jarlib=cdh4** o For MapR add: **jarlib=apache1.1.2**
kerbrealm	MapReduce	<none>	If the Hadoop cluster is Kerberos secured, use this value to specify the Kerberos Realm. For example, **kerbrealm=TIDALSOFT.LOCAL**
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**

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. Defaults to 50 if not specified.
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 in service.props of this adapter.
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 in the service.props of this adapter.
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 *Informatica Adapter Guide* for how to set this for Windows and Unix environments.
TDLJDBC_LIBPATH	JDBC (Windows only, optional)	<none>	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.
TDLINFA_REQUESTTIMEOUT	Informatica	<none>	(Optional) – The number of seconds before an API request times out. The default is 120 seconds, if not specified.
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.