



Cisco Workload Automation Amazon EC2 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 Workload Automation Amazon EC2 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	3
Introducing the Amazon EC2 Adapter	1-5
Overview	1-5
Instance Types	1-6
Instance Storage	1-6
Amazon Machine Image (AMI)	1-6
Regions and Availability Zones	1-6
Amazon Elastic Block Store (EBS)	1-7
Amazon Simple Storage Service (S3)	1-7
Configuring the Amazon EC2 Adapter	2-9
Overview	2-9
Licensing an Adapter	2-9
Securing the Adapter	2-10
About Amazon EC2 Credentials	2-10
Defining Runtime Users	2-11
Authorizing Schedulers to Work With EC2 Jobs	2-12
Defining a Security Policy for Amazon EC2	2-12
Defining a Scheduler User	2-13
Defining a Connection	2-15
Adding a Connection	2-15
Using the Amazon EC2 Adapter	3-19
Defining Amazon EC2 Jobs	3-19
Defining an Amazon EC2 Job	3-19
Monitoring Amazon EC2 Job Activity	3-24
Controlling Adapter and Agent Jobs	3-26
Holding a Job	3-26

Aborting a Job 3-26
Rerunning a Job 3-27
Making One Time Changes to an Adapter or Agent Job Instance 3-27
Deleting a Job Instance before It Has Run 3-27

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 Amazon EC2 Adapter with Cisco Workload Automation (CWA).

Audience

This guide is for administrators who install and configure the Amazon EC2 Adapter for use with Cisco Workload Automation, and who troubleshoot CWA installation and requirements issues.

Related Documentation

See the *Cisco Workload Automation 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 Cisco Workload Automation 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 Amazon EC2 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>Cisco Tidal Enterprise Scheduler Installation Guide</i>; usage provided in online Help only.
6.2.1 SP3	May 2016	<ul style="list-style-type: none">Consolidated all Amazon EC2 Adapter documentation into one document.



Introducing the Amazon EC2 Adapter

This chapter provides an overview of the Cisco Workload Automation Amazon EC2 (Elastic Compute Cloud) Adapter and its requirements:

- [Overview](#)
- [Instance Types](#)
- [Instance Storage](#)
- [Amazon Machine Image \(AMI\)](#)
- [Regions and Availability Zones](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)

Overview

Amazon EC2 is an Amazon Web service that provides scalable compute-capacity in the cloud. Specifically, EC2 provides the ability to run virtual machine (VM) instances on Amazon's cloud infrastructure. EC2 is elastic, which means instances can be spun up or down on the fly. EC2 is also configurable. You can choose from a large set of VM images that come pre-installed with the latest operating systems and applications. Or, you can upload your own custom VM images to EC2. You can also choose the exact CPU, memory, and storage configurations to suit your computing needs.

The Amazon EC2 Adapter, in conjunction with Cisco Workload Automation, can be used to define, launch, control, and monitor jobs submitted to EC2 on a scheduled basis. The Adapter integrates seamlessly in an enterprise scheduling environment.

The Amazon EC2 Adapter includes the following features:

- Connection management to monitor system status with a live connection to the EC2 Server.
- EC2 job management includes the following:
 - Scheduling and monitoring of jobs from a centralized work console with Cisco Workload Automation.
 - Defined dependencies and events with Cisco Workload Automation for scheduling control.

Instance Types

There are several instance types to choose from, each with different configurations of memory and CPU capacity.

- **Standard Instances** – Well suited for most applications.
- **Micro Instances** – Provide a small amount of consistent CPU resources and allows applications to increase CPU capacity in short bursts when additional cycles are available.
- **High-memory Instances** – Offer large memory sizes for high throughput applications, including database and memory caching applications.
- **High-CPU Instances** – Has more CPU capacity than memory and well suited for compute-intensive instances.
- **Cluster Compute Instances** – Provides proportionally high CPU resources with increased network performance and are well suited for High Performance Compute (HPC) applications and other demanding network-bound applications.
- **Cluster GPU Instances** – Provide general-purpose graphics processing units (GPUs) with proportionally high CPU and increased network performance for applications benefitting from highly parallelized processing, including HPC, rendering and media processing applications.

Instance Storage

There are two types of instance storage:

- **Instance storage** – keeps the instance boot partition on a temporary store that is destroyed when the instance lifecycle ends.
- **EBS storage** – keeps the instance's boot partition on an EBS volume and can be persisted beyond the lifecycle of the instance.

Amazon Machine Image (AMI)

Amazon Machine Images (AMIs) are VM images from which instances are launched. They can be installed with a number of popular operating systems and applications. EC2 has a large number of existing AMIs to choose from. Or, customers can create their own AMIs from a VM image they upload themselves into EC2.

Regions and Availability Zones

Amazon segments cloud infrastructure into regions and availability zones. Regions are segmented geographically. There are major regions for both coasts of the United States for example. Regions are further segmented into availability zones. Each availability zone is a logically distinct “data center” that is insulated from other availability zones. By keeping instances in separate availability zones, customers can protect their applications from failures in other zones. EC2 provides inexpensive, low latency network connectivity between availability zones in the same region.

Amazon Elastic Block Store (EBS)

Amazon Elastic Block Store (EBS) is one of the persistent storage types for EC2. The other is Amazon Simple Storage Service (S3). EBS stores or volumes act as raw unformatted external hard drives that can be formatted using a file system such as ext3 (Linux) or NTFS (Windows) and mounted on an EC2 instance.

If an EBS volume is used as the boot partition of an instance, that instance can be stopped and restarted. In contrast, instances cannot be restarted after being stopped if their boot partition is stored in instance storage.

Snapshots can be created of EBS volumes and backed up into S3. These snapshots can be restored back into an EBS volume at any time.

Amazon Simple Storage Service (S3)

S3 is the other persistent storage type for EC2. It provides a simple web service interface that customers use to store and retrieve any amount of data. Like a file system, S3 data is conceptualized into buckets (folders) and objects (files).



Configuring the Amazon EC2 Adapter

Overview

The Amazon EC2 Adapter software is installed as part of a standard installation of Cisco Workload Automation. However, you must perform the following steps to license and configure the Adapter before you can schedule and run EC2 jobs:

- [Licensing an Adapter](#) – Apply the license to the Amazon EC2 Adapter. You cannot define a EC2 connection until you have applied the EC2 license.
- [Securing the Adapter](#) – Define EC2 users that the Adapter can use to establish authenticated sessions with the EC2 server and permit requests to be made on behalf of the authenticated account.
- [Defining a Connection](#) – Define an EC2 connection so the master can communicate with the EC2 server.

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

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 Amazon EC2 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 EC2 jobs represent those users and passwords required for authentication. EC2 operations require authentication against a valid EC2 user as defined by a EC2 administrator.
- **Schedulers**
Schedulers are those users who will define and/or manage EC2 jobs. There are three aspects of a user profile that grant and/or limit access to scheduling jobs that affect EC2:
 - Security policy that grants or denies add, edit, delete and view capabilities for EC2 jobs.
 - Authorized runtime user list that grants or denies access to specific authentication accounts for use with EC2 jobs.
 - Authorized agent list that grants or denies access to specific Amazon EC2 Adapter connections for use when defining EC2 jobs.

About Amazon EC2 Credentials

Access to Amazon EC2 is granted using a certificate and access/secret key pair.

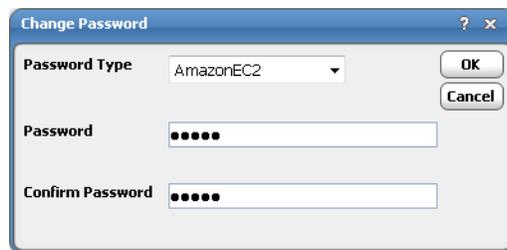
The access and secret keys are unique strings assigned to each customer where their account is created. Both are needed to authenticate into the EC2 system. The access key is considered public, while only the customer should know the secret key.

The Amazon EC2 adapter needs both keys to make a connection to EC2. Since each key-pair uniquely identifies a particular customer, they can be treated as a user name and password. As such, Cisco Workload Automation users can be used to store the access and private keys. When defining an EC2 connection, the user must select an existing EC2 user with the desired keys.

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 AWS access key ID in the **User Name** field.
- Step 4** For documentation, enter the **Full Name** or description associated with this user.
- Step 5** To define this user as a runtime user for EC2 jobs, click **Add** on the **Passwords** tab.
The **Change Password** dialog displays.



- Step 6** Select **AmazonEC2** from the **Password Type** list.
- Step 7** Enter a password (along with confirmation) in the **Password/Confirm Password** fields.
Only those users with a password specified for EC2 will be available for use with EC2 jobs. The password might be the same as the one specified for Windows/FTP jobs.
- Step 8** Click **OK** to return to the **User Definition** dialog.
The new password record displays on the **Passwords** tab.

The screenshot shows the 'User Definition' dialog box with the following fields and controls:

- User Name:** AKIAJ4C7NOYSARBJWRBQ
- Full Name:** (empty)
- Domain:** (empty)
- Windows/FTP/DataMover:** (empty)
- Adapter/Password Table:**

Adapter	Password	Description
AmazonEC2 Password	***	
- Buttons:** OK, Cancel, Add, Edit, Delete.

9 Click **OK** to add or save the user record in the Cisco Workload Automation database.

For further information about the **User Definition** dialog, see the “Users” chapter in your *Cisco Workload Automation User Guide*.

Authorizing Schedulers to Work With EC2 Jobs

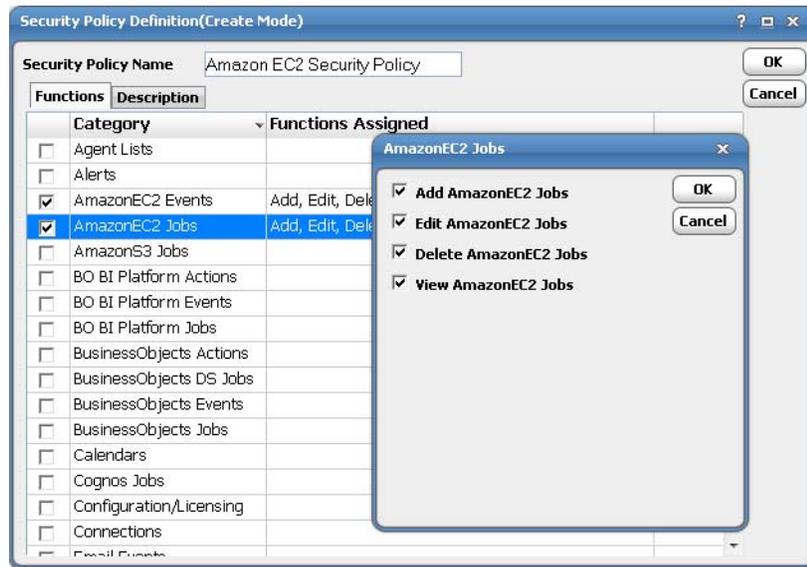
There are two steps for authorizing schedulers to work with Amazon EC2 jobs:

- [Defining a Security Policy for Amazon EC2](#)
- [Defining a Scheduler User](#)

Defining a Security Policy for Amazon EC2

To define a security policy for Amazon EC2:

-
- 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**.



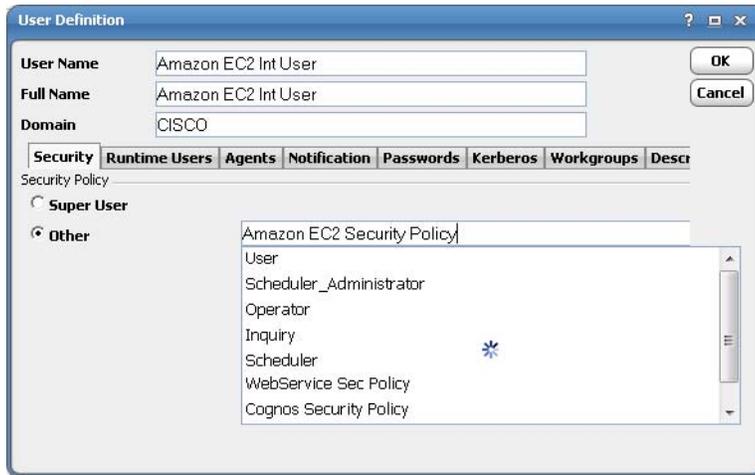
- Step 3** In the **Security Policy Name** field, enter a name for the policy.
- Step 4** On the **Functions** page, scroll to the **EC2 Jobs** category, click the ellipses on the right-hand side of the dialog and select the check boxes next to the functions that are to be authorized under this policy (**Add**, **Edit**, **Delete** and **View EC2s Jobs**).
- Step 5** Click **Close** on the **Function** drop-down list.
- Step 6** Click **OK** to save the policy.

For further information about setting up security policies, see your *Cisco Workload Automation User Guide*.

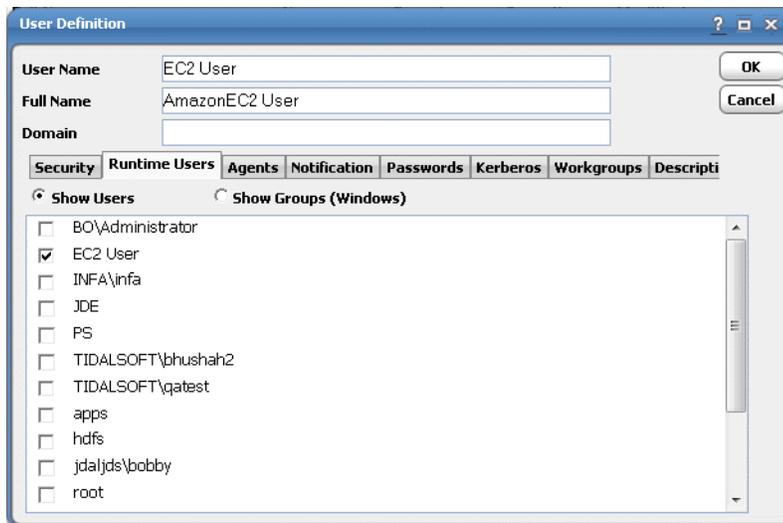
Defining a Scheduler User

To define a Scheduler user to work with EC2 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 EC2 jobs.
- Step 7** Click the **Runtime Users** tab.



- Step 8** Select the Amazon EC2 users that this scheduling user may use for AmazonEC2 authentication in AmazonEC2 jobs.
- Step 9** Click the **Agents** tab.
- Step 10** Select the check boxes for the EC2 connections that this scheduling user can access when scheduling jobs.
- Step 11** Click **OK** to save the user definition.

Defining a Connection

You must create one or more EC2 connections before Cisco Workload Automation can run your EC2 jobs. These connections also must be licensed before Cisco Workload Automation can use them. A connection is created using the **Connection Definition** dialog.

An EC2 connection represents a connection to an EC2 availability zone. All jobs executing under an EC2 connection would thus operate only on the instances in the connection's availability zone.

An EC2 connection needs three pieces of information: the availability zone (endpoint), the access key and the secret key. The zone can be selected from a dropdown list of zones. The dropdown list will be populated via a query to EC2. The access key and secret key can be selected indirectly by a dropdown list of EC2 users. As mentioned each the EC2 user's user name and password map to the access and secret keys, respectively.

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 > AmazonEC2 Adapter** from the context menu.

The **Amazon EC2 Adapter Connection Definition** dialog displays.

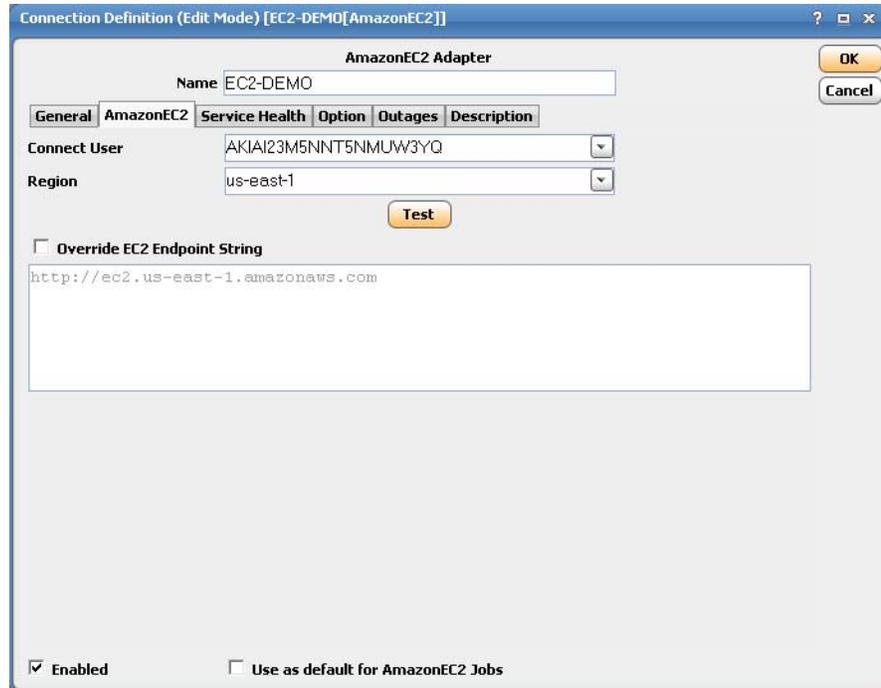
- Step 3** On the **General** page, enter a name for the new connection in the **Name** field.

Step 4 In the **Job Limit** field, select the maximum number of concurrent active processes that Cisco Workload Automation should submit to the Amazon EC2 server at one time.

Step 5 From the **Default Runtime User** drop-down list, you have the option to select the name of a default user for AmazonEC2 jobs.

Only authorized users that have been defined with Amazon EC2 passwords display in this list. The selected user is automatically supplied as the default runtime user in a new Cisco Workload Automation Amazon EC2 job definition.

Step 6 Click the **AmazonEC2** tab.



Step 7 From the **Connect User** list, select the associated Runtime User for Amazon EC2 to be used for connecting to Amazon EC2.

This is a persistent user connection to AmazonEC2 that is only used for administration and monitoring and for jobs with a matching runtime user. Jobs with a different runtime user specified will create additional temporary connections.

Step 8 From the **Region** list, enter the name of your Amazon EC2 server.

Step 9 In the **Override EC2 Endpoint String** field, enter the URL of the EC2 server in the format specified above.

Step 10 Click the **Test** button to test connectivity using the specified EC2 Server URL.

If the test is successful, the following information dialog box displays:



- Step 11** Click the **Service Health** tab to display the status and additional information of the for the selected availability zone.

Connection Definition (Edit Mode) [EC2-DEMO[AmazonEC2]]

AmazonEC2 Adapter

Name: EC2-DEMO

General AmazonEC2 Service Health Option Outages Description

Service Status

Availability Zone Status

Current Status	Details	
us-east-1a	available	
us-east-1b	available	
us-east-1c	available	
us-east-1d	available	

Refresh

Service Health

Enabled Use as default for AmazonEC2 Jobs

- Step 12** Click the **Options** tab to specify additional connection parameters. Setting these parameter values at the connection level will globally set the Job Server 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 job definition.



Note As of Cisco Workload Automation (Tidal Enterprise Scheduler) version 6.1, connection parameters for the Amazon EC2 Adapter is not supported.

- Step 13** Click **OK** to save the new EC2 connection.

The status light next to the connection indicates whether the Cisco Workload Automation Master is connected to the EC2 server. If the light is green, the EC2 server is connected.

A red light indicates that the master cannot connect to the EC2 server. EC2 jobs will not be submitted without a connection to the EC2 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 EC2 Connection** dialog describing the problem. Or, go to **Operator>Logs** to look for error messages associated with this connection.



Using the Amazon EC2 Adapter

This chapter covers these topics:

- [Defining Amazon EC2 Jobs](#)
- [Monitoring Amazon EC2 Job Activity](#)
- [Controlling Adapter and Agent Jobs](#)

Defining Amazon EC2 Jobs

EC2 Adapter jobs perform several useful tasks that would otherwise have to be manually initiated through the EC2 Web console or a command line tool. As with other Cisco Workload Automation adapters, Cisco Workload Automation EC2 jobs can be scheduled, controlled and coordinated within the Cisco Workload Automation ecosystem. The EC2 Adapter currently supports 5 job types, including three for starting, stopping and terminating EC2 instances and two for snapshotting EBS volumes and moving user-owned AMIs between different EC2 regions.

Defining an Amazon EC2 Job

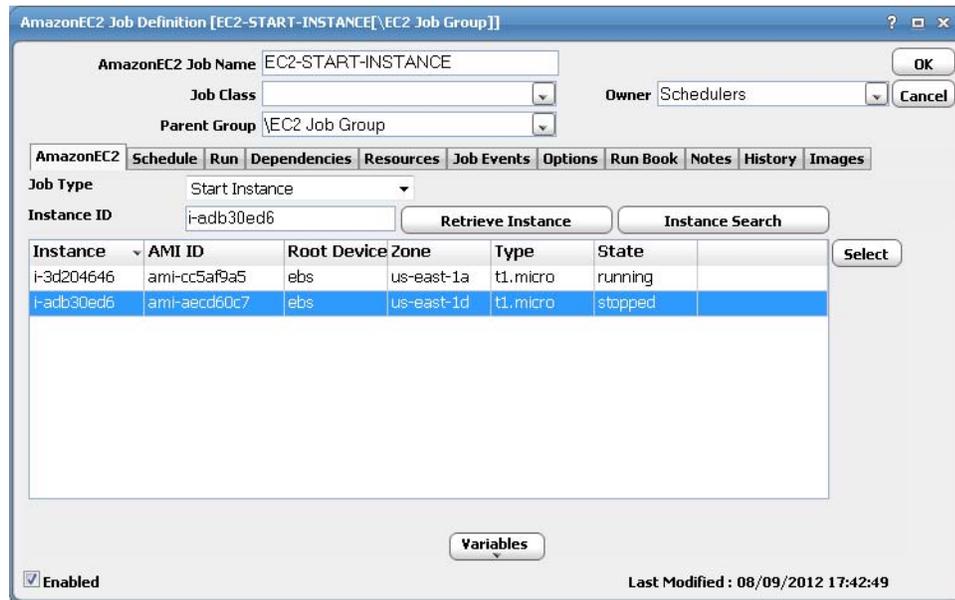
To define a Amazon EC2 job:

-
- Step 1** In the **Navigator** pane, select **Definitions>Jobs** to display the **Jobs** pane.
- Step 2** Right-click **Jobs** and select **Add>Amazon EC2 Job** from the context menu. The **Amazon EC2 Job Definition** dialog displays.

The screenshot shows the 'AmazonEC2 Job Definition' dialog box with the 'Run' tab selected. The 'Job Name' is 'EC2-SNAPSHOT-EBS-VOLUMES', 'Job Class' is empty, 'Parent Group' is 'EC2 Job Group', and 'Owner' is 'qatest'. In the 'Agent/Adapter Information' section, 'Agent/Adapter Name' is 'Amazon EC2[AmazonEC2]', 'Agent List' is empty, and 'Runtime User' is 'AKIAJ4C7NOYSARBJWRBQ'. The 'Tracking' section has 'Use' set to 'Exit code'. The 'Duration' section has 'Estimated', 'Minimum', and 'Maximum' all set to '1:00'. The 'Exclude Completed Abnormally' checkbox is checked. The job is 'Enabled' and was last modified on '07/24/2012 16:53:09'.

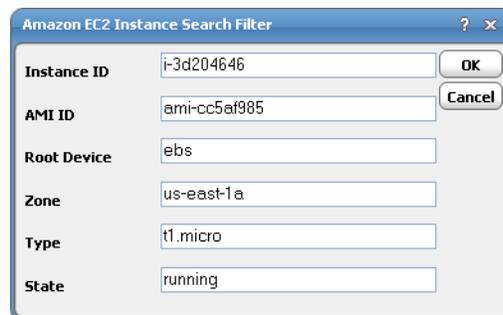
The **Run** tab is selected by default. You must first specify a name for the job, the Amazon EC2 Adapter connection that will be used for the job and a valid runtime user who has the appropriate authority for the batch job being scheduled.

- 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 drop-down list. This field is optional.
 - **Owner** – Select the Amazon EC2 job owner.
 - **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 Amazon EC2 Adapter connection to be used for this job from the drop-down list.
 - **Agent List Name** – Select a list for broadcasting the job to multiple Amazon EC2 servers rather than a specific agent/adaptor.
 - **Runtime User** – Select a valid runtime user with the appropriate authority for the job from the drop-down list.
- Step 5** Specify the appropriate **Tracking** and **Duration** information for the job. Refer to the *Cisco Workload Automation User Guide* for information on these options.
- Step 6** Click the **AmazonEC2** tab.



Step 7 From the **Job Type** list, select the job type. The following types are available:

- **EC2-START-INSTANCE** – Select this type to start an EC2 instance. Only one instance can be started at a time. Click **Retrieve Instance** to retrieve all EC2 instances. Select one of the retrieved instances and click **Select Instance**. Additionally, you can search for an instance via the **Amazon EC2 Instance Search Filter** dialog by clicking **Instance Search**.



- **EC2-STOP-INSTANCE** – Select this type to stop an EC2 instance. Only one instance can be stopped at a time. If the instance is store backed, the instance lifecycle ends. If the instance is an EBS-stored backed instance, the instance can be restarted later. Click **Retrieve Instance** to retrieve all EC2 instances. Select one of the retrieved instances and click **Select Instance**. Additionally, you can search for an instance via the **Amazon EC2 Instance Search Filter** dialog by clicking **Instance Search**.

AmazonEC2 Job Definition [EC2-STOP-INSTANCE[\EC2 Job Group]]

AmazonEC2 Job Name: EC2-STOP-INSTANCE

Job Class: []

Parent Group: \EC2 Job Group

Owner: Schedulers

AmazonEC2 Schedule Run Dependencies Resources Job Events Options Run Book Notes History Images

Job Type: Stop Instance

Instance ID: i-adb30ed6

Instance	AMI ID	Root Device	Device Zone	Type	State	Select
i-3d204646	ami-cc5af9a5	ebs	us-east-1a	t1.micro	running	
i-adb30ed6	ami-aecd60c7	ebs	us-east-1d	t1.micro	stopped	

Variables

Enabled Last Modified : 08/09/2012 17:40:30

- **EC2-TERMINATE-INSTANCE** – Select this type to terminate a EC2 instance that is running. Only one instance can be terminated at a time. Click **Retrieve Instance** to retrieve all EC2 instances. Select one of the retrieved instances and click **Select Instance**. Additionally, you can search for an instance via the **Amazon EC2 Instance Search Filter** dialog by clicking **Instance Search**.

AmazonEC2 Job Definition [EC2-TERMINATE-INSTANCE[\EC2 Job Group]]

AmazonEC2 Job Name: EC2-TERMINATE-INSTANCE

Job Class: []

Parent Group: \EC2 Job Group

Owner: qatest

AmazonEC2 Schedule Run Dependencies Resources Job Events Options Run Book Notes History Images

Job Type: Terminate Instance

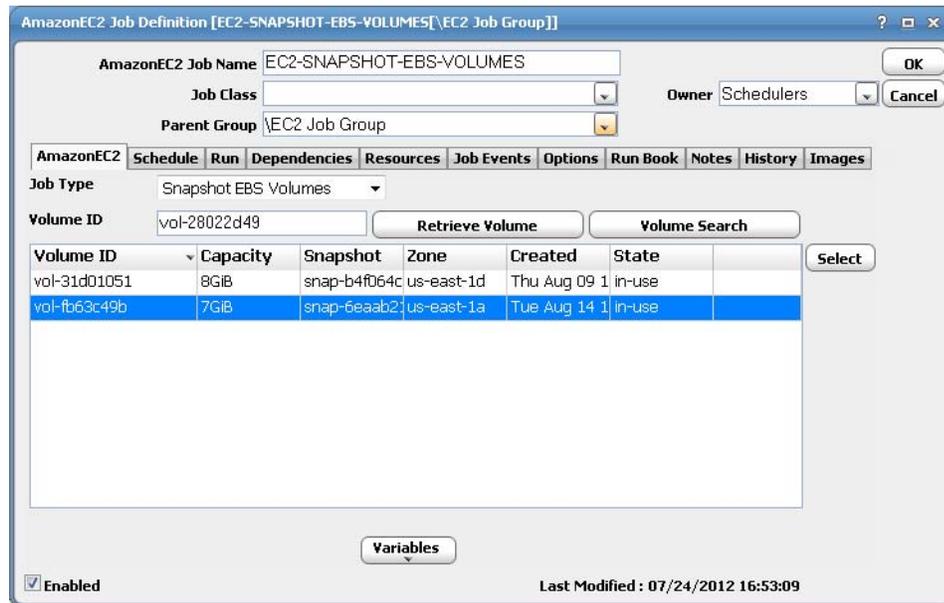
Instance ID: i-2667555e

Instance	AMI ID	Root Device	Device Zone	Type	State	Select
i-42335e38	ami-aecd60c7	ebs	us-east-1a	t1.micro	stopped	
i-e6335e9c	ami-aecd60c7	ebs	us-east-1a	t1.micro	running	

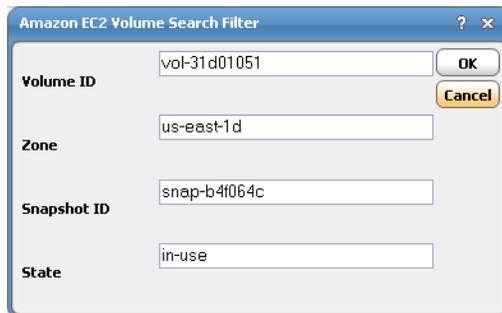
Variables

Enabled Last Modified : 07/24/2012 16:53:48

- **EC2-SNAPSHOT-EBS-VOLUMES** – Select this type to snapshot an existing EBS volume into S3. Only one EBS volume can be snapshot at a time. Click **Retrieve Volume** to retrieve all EC2 EBS volumes and click **Select**.



Additionally, you can search for an instance via the **Amazon EC2 Volume Search Filter** dialog by clicking **Volume Search**.



- **EC2-COPYAMI** – Copies an existing AMI in one EC2 region to another EC2 region. The connection user must own the AMI. Only one copy operation can be performed at a time. The next two steps describe how to define a Copy AMI job.

Step 8 Copy AMI from one region to another by following the steps below.



Note Before you can define a Copy AMI job type in Cisco Workload Automation (the next step), you must complete these prerequisite steps to copy AMI from one region to another.

- Upload your Amazon EC2 key.

```
$scp -I <keypair_name> <private_keyfile> <certificate_file> <username>@<dns_location>
<instance store directory>
```

- Bundle a customized image (requires root privileges).

```
# ec2-bundle-vol -k <private_keyfile> -c <certificate_file> -u <user_id>
```

- Migrate the manifest file to the destination region to verify usage in the destination.

```
ec2-migrate-manifest -m <manifest file> -c <certificate_file> -k <private_keyfile> -a
<access_key> -s <secret key> --region <destination region>
```

- d. Create a bucket in the source region.
- e. Upload a bundled AMI to the source bucket in the source region.

```
ec2-upload-bundle -b <bucket_path> -m <manifest_path> -a <access_key> -s <secret_key>
```

Step 9 Define a Copy AMI job type in Cisco Workload Automation:

- a. From the **Job Type** list, select **Copy AMI**.

- b. From the **Destination Region** list, specify the region you want to copy the AMI to.
- c. In the **AMI ID** field, enter the unique ID assigned to the AMI.
- d. In the **Source Bucket** field, enter the source bucket name you want to copy the AMI from.
- e. In the **Destination Bucket** field, enter the destination bucket name you want to copy the AMI to.
- f. In the **Source Manifest File** field, enter the path for the source AMI manifest.

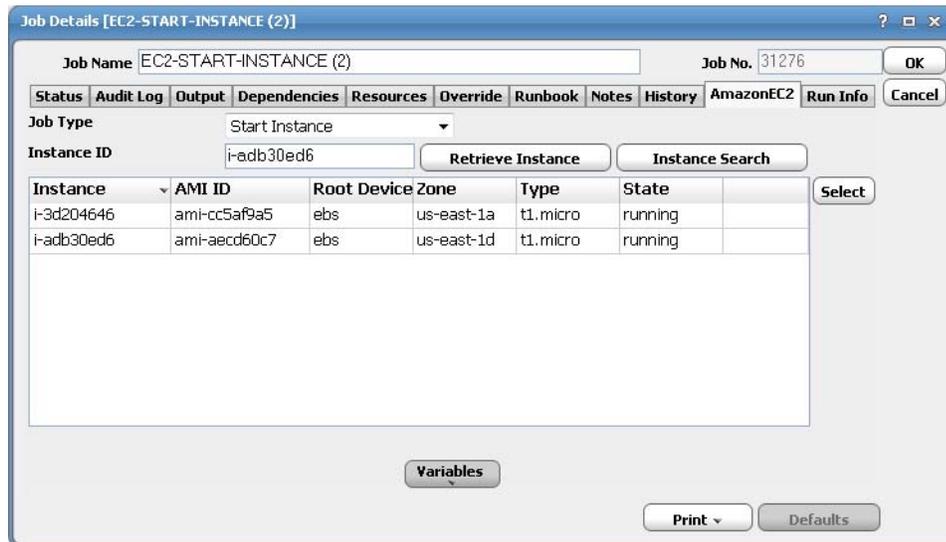
Step 10 Click **OK** to save the job.

Monitoring Amazon EC2 Job Activity

As Amazon EC2 job types run as pre-scheduled jobs, you can monitor the jobs as you would any other type of job in Cisco Workload Automation 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 Workload Automation User Guide* for instructions on using **Business Views**).

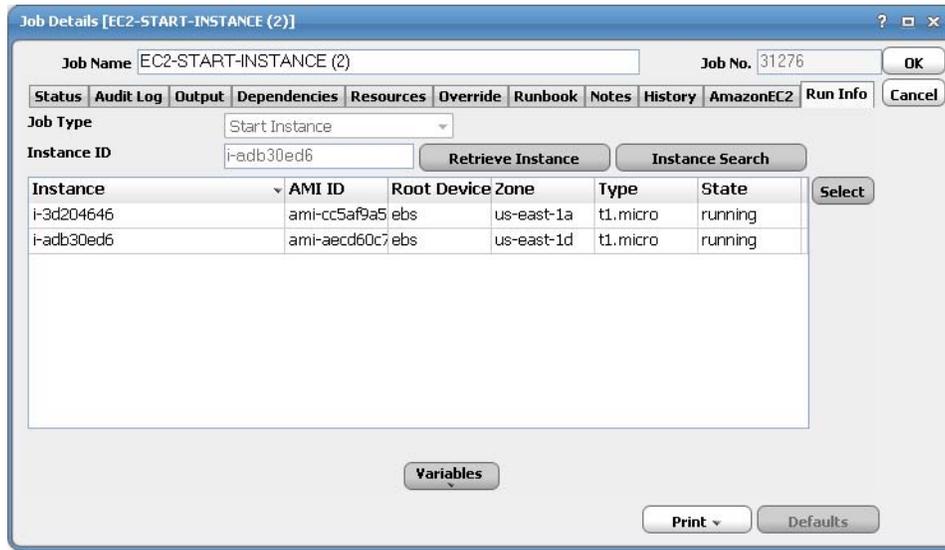
To monitor job activity:

- 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 job.
- Step 3** Click the **Output** tab to view the job output.
- Cisco Workload Automation 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.
- Step 4** Click the **AmazonEC3** tab to view the original request along with variables used when this job was submitted. This tab allows you to override the output format and parameter values prior to run or rerun. Overrides are not permitted when the job is running.



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 instance in the Amazon EC2 job. This tab is read-only.



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.