

NetIQ Cloud Manager 2.3

Procedures Guide

December 11, 2014

Legal Notice

THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT ARE FURNISHED UNDER AND ARE SUBJECT TO THE TERMS OF A LICENSE AGREEMENT OR A NON-DISCLOSURE AGREEMENT. EXCEPT AS EXPRESSLY SET FORTH IN SUCH LICENSE AGREEMENT OR NON-DISCLOSURE AGREEMENT, NETIQ CORPORATION PROVIDES THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW DISCLAIMERS OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS; THEREFORE, THIS STATEMENT MAY NOT APPLY TO YOU.

For purposes of clarity, any module, adapter or other similar material ("Module") is licensed under the terms and conditions of the End User License Agreement for the applicable version of the NetIQ product or software to which it relates or interoperates with, and by accessing, copying or using a Module you agree to be bound by such terms. If you do not agree to the terms of the End User License Agreement you are not authorized to use, access or copy a Module and you must destroy all copies of the Module and contact NetIQ for further instructions.

This document and the software described in this document may not be lent, sold, or given away without the prior written permission of NetIQ Corporation, except as otherwise permitted by law. Except as expressly set forth in such license agreement or non-disclosure agreement, no part of this document or the software described in this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, or otherwise, without the prior written consent of NetIQ Corporation. Some companies, names, and data in this document are used for illustration purposes and may not represent real companies, individuals, or data.

This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. NetIQ Corporation may make improvements in or changes to the software described in this document at any time.

U.S. Government Restricted Rights: If the software and documentation are being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), in accordance with 48 C.F.R. 227.7202-4 (for Department of Defense (DOD) acquisitions) and 48 C.F.R. 2.101 and 12.212 (for non-DOD acquisitions), the government's rights in the software and documentation, including its rights to use, modify, reproduce, release, perform, display or disclose the software or documentation, will be subject in all respects to the commercial license rights and restrictions provided in the license agreement.

© 2014 NetIQ Corporation and its affiliates. All Rights Reserved.

For information about NetIQ trademarks, see https://www.netiq.com/company/legal/.

Contents

About NetIQ Corporation **9**
About this Book and the Library **11**

Part I Administering the Cloud Manager Orchestration Server and Agent **13**

1 Understanding Basic Orchestration Functions **15**

1.1 Observing Discovery Jobs . . . 15
1.2 Deploying a Sample Job . . . 17
1.3 Creating a User Account . . . 19
1.3.1 Opening the Users Monitor . . . 19
1.3.2 Automatically Registering a User . . . 20
1.3.3 Manually Registering a User . . . 21
1.3.4 Logging In a User for Manual Registration . . . 22
1.3.5 Directory Service Authentication (Optional) . . . 23
1.4 Running a Sample Job . . . 23
1.5 Looking at the Job After It Has Run . . . 24
1.5.1 Verification at the Command Line . . . 24
1.5.2 Verification at the Jobs Monitor . . . 25
1.6 Using the zosadmin Command to Gather Information . . . 26
1.7 Stopping and Starting Orchestration Components . . . 27
1.7.1 Stopping and Starting the Orchestration Server . . . 27
1.7.2 Stopping and Starting the Orchestration Agent . . . 28
1.7.3 Starting and Stopping the Orchestration Server Console . . . 28
1.8 Configuring the Orchestration Server for LDAP or ADS Authentication . . . 29
1.9 Viewing Orchestration Server Data in the Sentinel Event Source Server . . . 29
1.10 Orchestration Server Log Levels Mapped to Sentinel Log Levels . . . 31

2 Determining the Version of an Orchestration Component **33**

3 Changing Orchestration Server Default Parameters and Values **35**

4 Changing Orchestration Agent Default Parameters and Values **37**

5 Caching Computed Facts in a Grid with Large Numbers of Resources **39**

5.1 Walkthrough: Scheduling a System Job . . . 39
5.1.1 Deploying a Sample System Job . . . 40
5.1.2 Creating a New Schedule for the Job . . . 41
5.1.3 Defining the New Schedule . . . 42
5.1.4 Activating the New Schedule . . . 46
5.1.5 Running the New Schedule Immediately . . . 46

6 Managing Virtual Machine Hosts **47**

6.1 Discovering VM Hosts and Repositories . . . 47
6.2 Discovering VM Images . . . 48
6.3 Resynchronizing the VM Host's State . . . 49

6.4 Shutting Down VM Hosts . . . 49
6.5 Restarting VM Hosts . . . 50
6.6 Understanding VM Host Failover . . . 51

7 Managing Virtual Machines 53

7.1 Provisioning a Virtual Machine . . . 53
7.2 Provisioning Actions and History . . . 54
7.2.1 What are Provisioning Actions? . . . 54
7.2.2 How Actions Are Displayed in the Orchestration Console . . . 54
7.2.3 Prioritized Provisioning Actions . . . 57
7.3 Managing a Running Virtual Machine . . . 57
7.3.1 Using the Right-Click Menu for Provisioning Actions . . . 57
7.3.2 Releasing a Virtual Machine from Usage . . . 61
7.3.3 Managing Virtual Machine Templates . . . 61
7.3.4 Managed Virtual Machine Actions . . . 62
7.4 Resynchronizing the State of All VMs . . . 63
7.5 Resynchronizing the State of All VMs of a Specific VM Host . . . 63
7.6 Shutting Down Multiple VMs . . . 64
7.7 Destroying and Deleting a Virtual Machine . . . 64

8 Managing VM Repositories 67

8.1 Provisioning a VM from a Local Repository . . . 67
8.2 The Datagrid Repository . . . 68
8.3 Provisioning a VM from a NAS Repository . . . 68
8.4 The SAN Repository . . . 68

Part II Administering the Cloud Manager Application Server and Console 69

9 Setting Up the Cloud Environment 71

9.1 Creating a Zone . . . 71
9.2 Creating Resource Groups Within the Zone . . . 72
9.2.1 Resource Group Requirements . . . 72
9.2.2 Shared and Dedicated Resource Groups . . . 72
9.2.3 Creating a Resource Group . . . 72
9.3 Creating Service Levels for Resource Groups . . . 73
9.4 Creating an Organization . . . 75
9.5 Creating an Organization’s Business Groups . . . 77
9.6 Creating Users and Groups . . . 78
9.6.1 Manually Creating Users . . . 79
9.6.2 Manually Creating User Groups . . . 80
9.6.3 Importing Users from LDAP . . . 82
9.6.4 Importing User Groups from LDAP . . . 84

10 Configuring the Cloud Manager System 87

10.1 Managing Zones . . . 87
10.1.1 Creating Zones . . . 87
10.1.2 Disabling Zones . . . 88
10.1.3 Enabling Zones . . . 88
10.1.4 Removing Zones . . . 89
10.2 Customizing the Capacity Thresholds and Data Refresh Interval . . . 89
10.3 Managing System Users, User Groups, and Roles . . . 90

10.4 Configuring Email Notifications . . . 90
10.5 Configuring Remote Console Access to Workloads . . . 91
10.5.1 Disabling Remote Console Access . . . 91
10.5.2 Setting Up the Built-In VNC Repeater . . . 92
10.5.3 Setting Up an External VNC Repeater . . . 92
10.5.4 Setting Up Direct Connections . . . 92
10.5.5 Enabling Repeater SSL Encryption . . . 93
10.5.6 Enabling File Uploading . . . 93
10.6 Configuring Auto Approval for Business Service Requests . . . 94
10.7 Customizing the Cost Update Schedule . . . 94
10.8 Customizing the Cloud Manager Console Interface . . . 95
10.9 Setting Up the Web Application Launcher (Launch Pad) . . . 96
10.9.1 Configuring the Widget . . . 96

11 Setting Up and Managing Users 101

11.1 Creating User Accounts . . . 101
11.1.1 Manually Creating System and Organization Users . . . 101
11.1.2 Manually Importing System Users from LDAP . . . 102
11.1.3 Manually Importing Organization Users from LDAP . . . 103
11.2 Enabling Automatic User Creation . . . 105
11.2.1 Setting Up Automatic User Creation for System Users . . . 105
11.2.2 Setting Up Automatic User Creation for Organization Users . . . 106
11.2.3 Automating Role Assignments to Automatically Imported Users . . . 106
11.3 Creating User Groups . . . 107
11.3.1 Manually Creating System and Organization User Groups . . . 107
11.3.2 Importing System User Groups from LDAP . . . 108
11.3.3 Importing Organization User Groups from LDAP . . . 109
11.4 Assigning Roles to Users and Groups . . . 111
11.5 Deleting Users . . . 112
11.6 Deleting User Groups . . . 112

12 Managing the Workload Template Catalog 113

12.1 Creating Workload Templates . . . 113
12.2 Assigning Workload Templates to Organizations and Business Groups . . . 116
12.2.1 Assigning Workload Templates to an Organization . . . 116
12.2.2 Assigning Workload Templates to a Business Group . . . 116
12.2.3 Removing Workload Template Assignments from an Organization . . . 117
12.2.4 Removing Workload Template Assignments from a Business Group . . . 117
12.3 Modifying Workload Templates . . . 118
12.4 Deleting Workload Templates . . . 118
12.5 Add-on Applications . . . 119
12.5.1 Creating an Add-on Application . . . 119
12.5.2 Configuring an Add-on Application . . . 120
12.6 Add-On Services . . . 121
12.6.1 Creating an Add-on Service . . . 122
12.6.2 Configuring an Add-on Service . . . 123
12.7 Legal Agreements . . . 124
12.7.1 Creating a Legal Agreement . . . 124
12.7.2 Editing a Legal Agreement . . . 125
12.8 Resource Pricing Packages . . . 127
12.8.1 Creating a Resource Pricing Package . . . 127
12.8.2 Configuring a Resource Pricing Package . . . 128

13 Managing Organizations **131**

13.1 Creating Organizations . . . 131
13.2 Customizing Organization Configuration Settings . . . 133
13.3 Creating Business Groups . . . 134
13.4 Assigning Resource Groups to Organizations and Business Groups . . . 136
13.4.1 Assigning Resource Groups to an Organization . . . 136
13.4.2 Assigning Resource Groups to a Business Group . . . 137
13.4.3 Removing Resource Group Assignments from an Organization . . . 138
13.4.4 Removing Resource Group Assignments from a Business Group . . . 138
13.5 Managing Organization Users, User Groups, and Roles . . . 138

14 Managing Resources **141**

14.1 Creating, Modifying, and Deleting Resource Groups . . . 141
14.1.1 Creating Resource Groups . . . 141
14.1.2 Modifying Resource Groups . . . 143
14.1.3 Deleting Resource Groups . . . 143
14.2 Creating, Modifying, and Deleting Service Levels . . . 144
14.2.1 Creating Service Levels . . . 144
14.2.2 Modifying Service Levels . . . 145
14.2.3 Deleting Service Levels . . . 146
14.2.4 Creating Service Level Objectives . . . 146
14.2.5 Modifying Service Level Objectives . . . 147
14.2.6 Deleting Service Level Objectives . . . 148
14.3 Assigning Service Levels to Resource Groups . . . 148
14.3.1 Assigning Service Levels . . . 148
14.3.2 Removing Service Levels . . . 148
14.4 Editing Network Configurations . . . 149
14.4.1 Editing the General Network Configuration . . . 149
14.4.2 Editing Network Associations . . . 150
14.4.3 Editing a Network's IP Pool . . . 151
14.5 Monitoring Resource Capacity . . . 154
14.5.1 Accessing the Capacity View . . . 155
14.5.2 Understanding the Capacity View . . . 155
14.5.3 Updating the Capacity Data . . . 158
14.5.4 Debugging Capacity Collection Issues . . . 158

15 Managing Business Services **161**

15.1 Requesting Business Services . . . 161
15.2 Managing Business Service Requests . . . 163
15.2.1 Submitting a Request . . . 164
15.2.2 Editing a Request . . . 164
15.2.3 Withdrawing a Request . . . 164
15.2.4 Deleting a Request . . . 165
15.3 Importing a Single Virtual Machine into an Existing or a New Business Service . . . 165
15.4 Importing Multiple Virtual Machines into Cloud Manager Business Services . . . 166
15.4.1 Bulk Importing Multiple Virtual Machines . . . 167
15.4.2 About the Bulk Import Spreadsheet . . . 168
15.5 Completing Pre-Build Configuration Tasks . . . 175
15.6 Completing Post-Build Configuration Tasks . . . 175
15.7 Managing Deployed Business Services . . . 176
15.7.1 Starting, Suspending, or Shutting Down a Workload . . . 176
15.7.2 Opening a Workload Console . . . 176
15.7.3 Adding Files to a Workload . . . 177
15.7.4 Delegating Ownership of a Business Service . . . 178
15.8 Changing Deployed Business Services . . . 179

15.8.1 Changing a Business Service's Details . . . 179
15.8.2 Reassigning a Business Service to a Different Business Group . . . 179
15.8.3 Adding a Workload . . . 180
15.8.4 Modifying a Workload . . . 181
15.8.5 Removing a Workload . . . 182
15.9 Extending Business Service Expiration Dates . . . 182
15.10 Displaying or Hiding Business Service Costs . . . 183
15.10.1 Configuring an Organization's Costs Setting . . . 183
15.10.2 Configuring a Business Group's Costs Setting . . . 183
15.10.3 Configuring a User's Costs Setting . . . 184
15.11 Understanding Business Service Contracts . . . 184
15.11.1 Restrictions Resulting from an Active Contract . . . 184
15.11.2 Contract Change Period . . . 184
15.11.3 Contract Renewals . . . 184

16 Managing Tasks in the Business Service Workflow 187

16.1 Displaying Administrator or Sponsor Tasks . . . 187
16.2 Claiming Tasks . . . 187
16.3 Approving and Denying Requests . . . 188
16.4 Completing Configuration Tasks . . . 188
16.4.1 Completing Pre-Build Configuration Tasks . . . 189
16.4.2 Completing Post-Build Configuration Tasks . . . 189
16.5 Completing Trigger Tasks . . . 190
16.5.1 Completing Reboot Trigger Tasks . . . 190

17 Generating Cloud Manager Reports 193

17.1 Report Descriptions . . . 193
17.2 Generating Reports . . . 194

18 Changing Application Server Default Parameters and Values 195

18.1 Increasing the Java Heap Space on the Application Server . . . 195

A Setting Up Cloud Manager to Log to a Sentinel Collector 197

B Integrating Cloud Manager with the Nixu IP Address Management System 199

B.1 Configuring Nixu NameSurfer for Cloud Manager Integration . . . 199
B.2 Configuring Cloud Manager for Nixu NameSurfer Integration . . . 201

C Using REST APIs to Customize Cloud Manager Behavior 205

C.1 Cloud Manager REST API Overview . . . 205
C.2 Using the REST APIs to Modify the Cloud Manager Workflow . . . 205
C.2.1 Workflow Callout Registration Concepts . . . 206
C.2.2 Skipping the PreConfig Approval Stage by Invoking REST Callouts . . . 207
C.2.3 REST URLs for Configuring and Registering Custom Callouts . . . 208
C.2.4 Overriding the Default Task Order . . . 212

D Documentation Updates 215

D.1 December 15, 2014 . . . 215
D.2 December 11, 2014 . . . 215
D.3 October 15, 2014 . . . 216

About NetIQ Corporation

NetIQ, an Attachmate business, is a global leader in systems and security management. With more than 12,000 customers in over 60 countries, NetIQ solutions maximize technology investments and enable IT process improvements to achieve measurable cost savings. The company's portfolio includes award-winning management products for IT Process Automation, Systems Management, Security Management, Configuration Audit and Control, Enterprise Administration, and Unified Communications Management. For more information, please visit www.netiq.com.

Contacting Sales Support

For questions about products, pricing, and capabilities, please contact your local partner. If you cannot contact your partner, please contact our Sales Support team.

Worldwide:	www.netiq.com/about_netiq/officelocations.asp
United States and Canada:	888-323-6768
Email:	info@netiq.com
Web Site:	www.netiq.com

Contacting Technical Support

For specific product issues, please contact our Technical Support team.

Worldwide:	www.netiq.com/Support/contactinfo.asp
North and South America:	1-713-418-5555
Europe, Middle East, and Africa:	+353 (0) 91-782 677
Email:	support@netiq.com
Web Site:	www.netiq.com/support

Contacting Documentation Support

Our goal is to provide documentation that meets your needs. If you have suggestions for improvements, click **Add Comment** at the bottom of any page in the HTML versions of the documentation posted at www.netiq.com/documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

Contacting the Online User Community

Qmunity, the NetIQ online community, is a collaborative network connecting you to your peers and NetIQ experts. By providing more immediate information, useful links to helpful resources, and access to NetIQ experts, Qmunity helps ensure you are mastering the knowledge you need to realize the full potential of IT investments upon which you rely. For more information, please visit http://community.netiq.com.

About this Book and the Library

The *Procedures Guide* provides step-by-step procedures for common tasks to be performed by users of the NetIQ Cloud Manager product, regardless of their role in the organization. The tasks are organized roughly sequentially; tasks found earlier in this guide are typically executed prior to tasks later in the guide during the initial deployment and implementation phase of the product deployment project.

The guide is organized in the following general sections:

- Part I, “Administering the Cloud Manager Orchestration Server and Agent,” on page 13
- Part II, “Administering the Cloud Manager Application Server and Console,” on page 69
- Appendix D, “Documentation Updates,” on page 215

These sections include the detail you need to get your Cloud Manager system up and running in a test environment first, and eventually in a production environment.

Intended Audience

This book provides information for individuals responsible for understanding administration concepts and implementing a secure, distributed administration model.

Other Information in the Library

The library provides the following information resources:

Product Overview

Provides information about the NetIQ Cloud Manager product features, functionality, and concepts.

Installation Guide

Provides detailed planning and installation information.

Procedures Guide

Provides step-by-step guidance for many administration tasks.

Reference Guide

Provides detailed reference information about tools and interfaces used by this product.

Administering the Cloud Manager Orchestration Server and Agent

This section includes information you will need to know as you administer the Cloud Manager Orchestration Server.

- Chapter 1, “Understanding Basic Orchestration Functions,” on page 15
- Chapter 2, “Determining the Version of an Orchestration Component,” on page 33
- Chapter 3, “Changing Orchestration Server Default Parameters and Values,” on page 35
- Chapter 4, “Changing Orchestration Agent Default Parameters and Values,” on page 37
- Chapter 5, “Caching Computed Facts in a Grid with Large Numbers of Resources,” on page 39
- Chapter 6, “Managing Virtual Machine Hosts,” on page 47
- Chapter 7, “Managing Virtual Machines,” on page 53
- Chapter 8, “Managing VM Repositories,” on page 67

1 Understanding Basic Orchestration Functions

After you install and configure the basic components of the NetIQ Cloud Manager Orchestration components, (that is, the Orchestration Server, the Orchestration Agent, and the Orchestration Clients, including the Orchestration Console), you will want to see them at work. The information in this section is organized sequentially (that is, in a “walkthrough” scenario) so that you can follow the process an administrator might use to begin applying Cloud Manager Orchestration capabilities in a production environment.

- Section 1.1, “Observing Discovery Jobs,” on page 15
- Section 1.2, “Deploying a Sample Job,” on page 17
- Section 1.3, “Creating a User Account,” on page 19
- Section 1.4, “Running a Sample Job,” on page 23
- Section 1.5, “Looking at the Job After It Has Run,” on page 24
- Section 1.6, “Using the zosadmin Command to Gather Information,” on page 26
- Section 1.7, “Stopping and Starting Orchestration Components,” on page 27
- Section 1.8, “Configuring the Orchestration Server for LDAP or ADS Authentication,” on page 29
- Section 1.9, “Viewing Orchestration Server Data in the Sentinel Event Source Server,” on page 29
- Section 1.10, “Orchestration Server Log Levels Mapped to Sentinel Log Levels,” on page 31

The first three subsections listed above are basic tasks you need to perform to make the Orchestration system perform at a basic level. The other sections include information to help you understand how Cloud Manager Orchestration can work in your production environment.

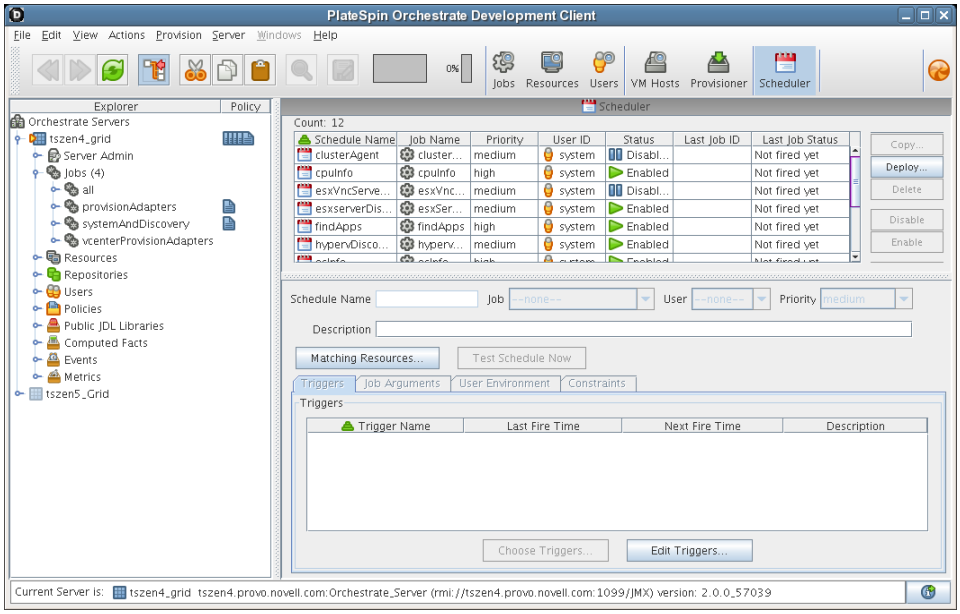
1.1 Observing Discovery Jobs

When you created a resource account for the first time, you might have noticed the status window of the Resource object change colors (see Step 2 in “Automatically Registering a Resource” in the *NetIQ Cloud Manager 2.3 Installation Guide*) from blue to green. You might also notice new jobs

displayed as objects in the Explorer panel. What you are observing are the “discovery” jobs that are shipped with the Orchestration Server (different discovery jobs are shipped with the Orchestration Server, depending on which management pack you license).

To understand the reason why these jobs run:

1 In the Orchestration Console, click *Scheduler* to open the Job Schedule view in the workspace.

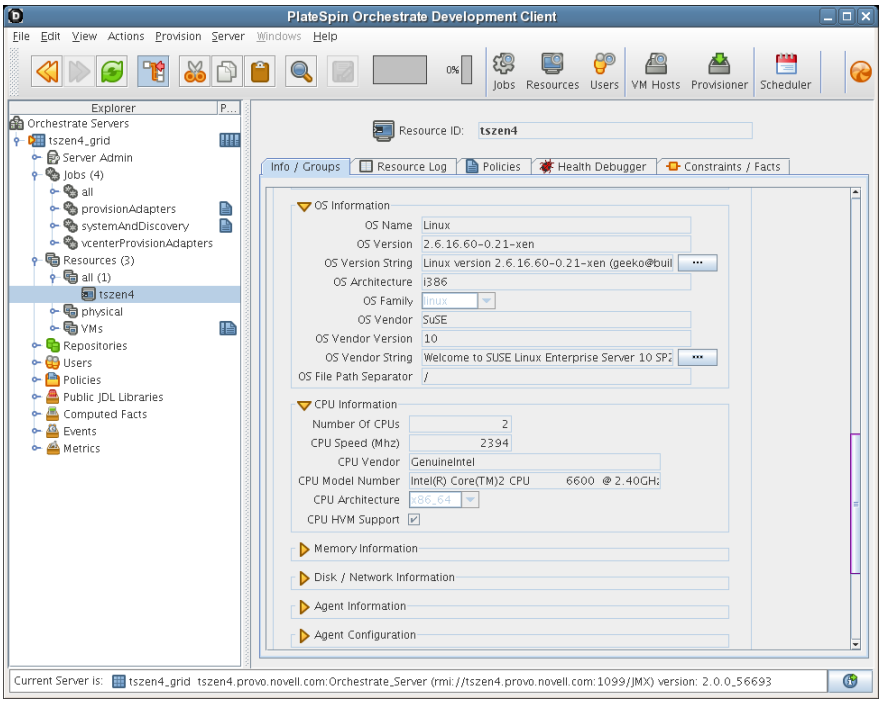


In this walkthrough of basic Orchestration Server functionality, you can see that several jobs are configured to run. If you select one of the jobs, such as `cpuInfo`, you will see that it was configured with a trigger called RESOURCE_ONLINE. All of the discovery jobs, like `cpuInfo`, are configured to run when the resource is online, that is, when the resource agent has logged into the Orchestration Server.

The discovery jobs, including provisioning adapter jobs, run basic operations at resource start as a convenience, to gather data that you or a job developer might need later when creating jobs, or that the Orchestration system might need as it allocates resources to run jobs. For example, the `cpuInfo` job and the `osInfo` job do some basic probing of the computing node (the machine where the agent is installed and has a resource account) for later reference.

To verify this, you can view the resource account that you created during the basic installation (as documented in “Creating a Resource Account” in the *NetIQ Cloud Manager 2.3 Installation Guide*) by selecting its object in the Explorer tree. By default, the *Info/Groups* page for the resource opens in the Orchestration Console admin view.

Figure 1-1 *Resource Information Page After Discovery Jobs Run*



If you scroll down on the *Info/Groups* page, you see that the discovery jobs have gathered basic data about the processor and operating system of this computing node. If the jobs had not run at resource start, this information about the resource would not be ready for use.

Now that you have seen a how jobs are run by the Orchestration system on resource start, you can walk through the process of deploying and running a sample job on your own by proceeding with Section 1.2, “Deploying a Sample Job,” on page 17.

1.2 Deploying a Sample Job

One of the main functions of the Orchestration Server is to run application requests, called jobs, on grid resources. Because the Orchestration Server is capable of handling multiple application requests, it uses a policy-based broker and scheduler to decide when and how a job should run on the resources. These decisions are based on many controlled factors, including the number of resource nodes, their cost, and a variety of other factors as requested by the application, but managed under policy constraints set up by the administrator or the job developer.

Developing a job involves the creation of an application executable and a job file.

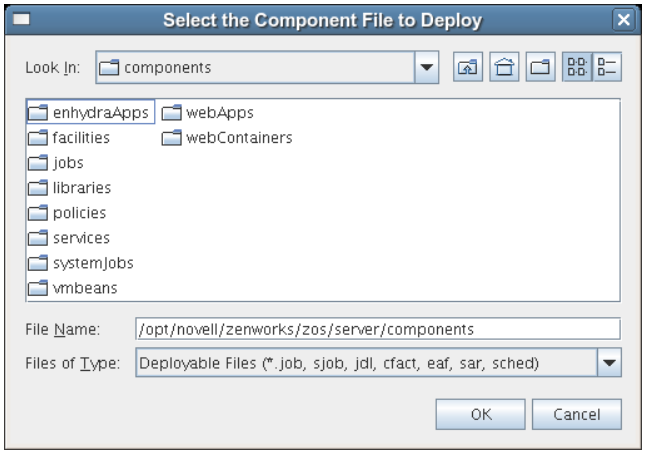
Before a job can run, the Orchestration Server administrator must deploy it, which involves moving it from a development state to a state where it is ready and available for users. Only the administrator has the necessary rights to deploy a job.

There are three methods you can use to deploy a job:

- Deploy from the Orchestration Console by right-clicking the *Jobs* container in the Explorer panel.
- Deploy from the Orchestration Console by selecting the *Actions* menu in the Orchestration Console.
- Deploy from the `zosadmin` command line (`zosadmin deploy` *`path_to_job`*).

For this walkthrough, we will deploy a simple job developed for Orchestration Server customers to demonstrate how jobs are deployed and run. Although the walkthrough shows only the first method for deploying, the other methods are relatively simple, so no further explanation is provided.

1. In the Explorer panel of the Orchestration Console, right-click the *Jobs* object, then click *Deploy Job* to open the Select the Component File to Deploy dialog box.

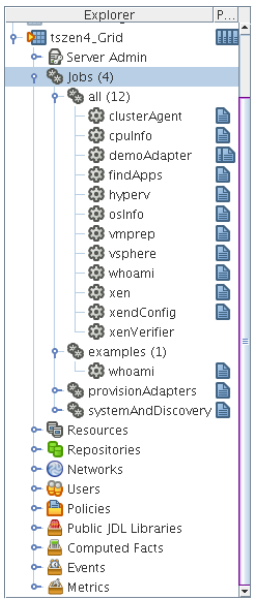


2. Open the *Look In* drop-down list, then navigate to the location of the job you want to deploy.

   Although a job developer can store Orchestration Server jobs at any location on the network, the sample jobs shipped with the Orchestration Server are limited to the directories where the product is installed. For this walkthrough, navigate to the `/opt/novell/zenworks/zos/server/examples` directory on the Orchestration Server. If the Orchestration Console is installed on a Windows machine, the default location is `c:\Program Files\Novell\zos\clients\examples`.

3. Select `whoami.job`, then click *OK* to deploy the job to the *Jobs* container.

   The *whoami* job appears in the *all* container and in the *examples* container in the tree.



   When deployed, the job is sent over the wire to the Orchestration Server with which it is associated. It is persisted there until undeployed.

When the job is available, you need to create a user who can run that job. For more information, see Section 1.3, “Creating a User Account,” on page 19

1.3 Creating a User Account

Although the Orchestration Server has some pre-assembled jobs, such as the `cpuInfo` discovery job that you learned about earlier, most jobs must be developed by a job developer, then be run and managed by a user (also called a job manager). Without an authorized individual who can log in to the Orchestration Server system to manage the use of a job, the product does not realize its potential.

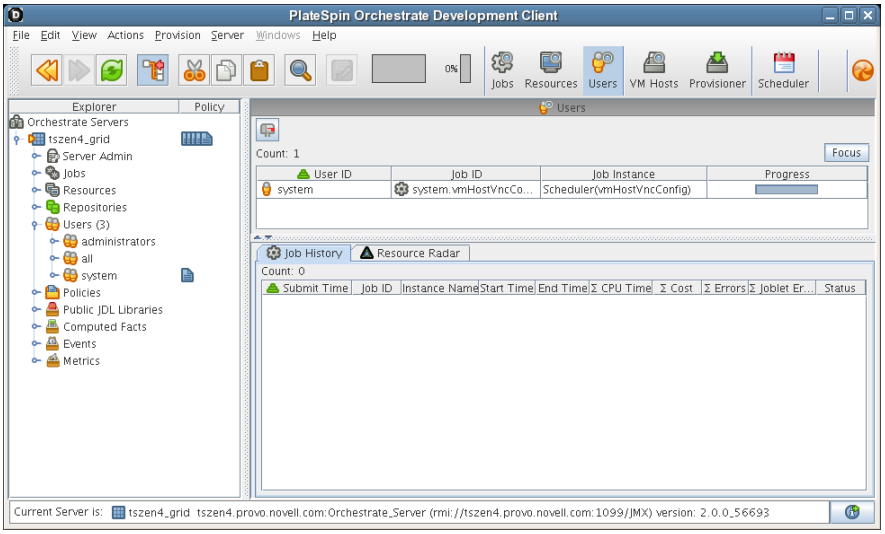
This section of the walkthrough introduces the basics of creating a user account:

- Section 1.3.1, “Opening the Users Monitor,” on page 19
- Section 1.3.2, “Automatically Registering a User,” on page 20
- Section 1.3.3, “Manually Registering a User,” on page 21
- Section 1.3.4, “Logging In a User for Manual Registration,” on page 22
- Section 1.3.5, “Directory Service Authentication (Optional),” on page 23

1.3.1 Opening the Users Monitor

Now that the Orchestration Server has run discovery jobs and you have deployed a sample job, you can begin to create user accounts. To do so, open the Orchestration Console and click *Users* in the toolbar to open the Users Monitor in the Workspace panel of the Orchestration Console.

Figure 1-2 *Users Monitor of the Orchestration Server Console*



In this monitor, you can see the users that are connected to the server and what they are doing in the grid.

If a user logs in but has not been registered (that is, no account is created for that user), the authentication to the server is retried every 90 seconds. If this is the case, the User Registration icon has a “flag up” status, meaning that a user is waiting to register. If the icon has a “flag down” status, either no user accounts have been created or all active users are logged in, so none are waiting to register.

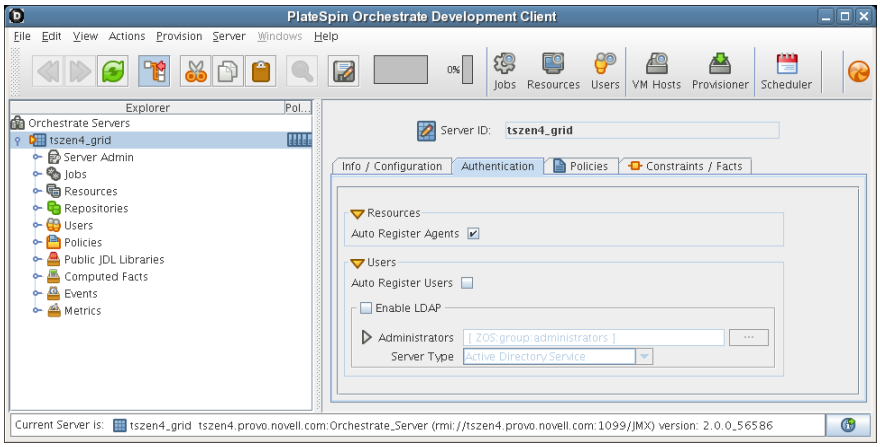
You can use the Orchestration Console to register a user automatically (see Section 1.3.2, “Automatically Registering a User,” on page 20) or to register a user manually (see Section 1.3.3, “Manually Registering a User,” on page 21). You can also select which users can log in to create accounts (see “Selecting a Resource for Manual Registration” in the *NetIQ Cloud Manager 2.3 Installation Guide*).

The Users Monitor has many features to help you manage users when they are registered, including the jobs and joblets assigned to individual users. For more detailed information about the Users Monitor, see Section 1.3.1, “Opening the Users Monitor,” on page 19.

1.3.2 Automatically Registering a User

If your network environment does not require a high level of security (such as in a development and testing environment) and you want a quick way to create a user account without a password, you can do so at the Orchestration Console.

1. In the Explorer tree of the Orchestration Console, select the grid object representing the Orchestration Server to open the *Info/Configuration* page of the grid object, then select the *Authentication* tab to open the Authentication page.
2. In the *Users* section of the page, select the *Auto Register Users* check box, then click the Save icon.



3. Use the `zos` command line interface to log in to the server.

   3a From a system terminal, enter the following command:

   ```
   zos login -u user_ID
   ```

   If you are attempting to log in to a machine other than the local host, you can alter the command to the following:

   ```
   zos login Orchestration_Server_name -u user_ID
   ```

   3b When prompted for the user password, press Enter.

   3c (Conditional) If you are prompted for a decision regarding whether you want to accept the server certificate, enter `yes`.

NOTE: You can assign a password for the user at a later time in the *Info/Groups* page of the User Object.

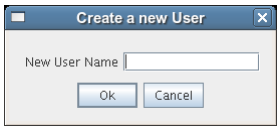
When a user logs out, the User object icon is dimmed in the Explorer tree or in the admin view of each User group to which it belongs.

1.3.3 Manually Registering a User

If you want a higher level of security for authorized users, you can manually create a user account in the Orchestration Console before the user logs in. When a user is created in the Orchestration Server Console, that user is ready to run jobs.

To create a new user in the Orchestration Console Explorer tree:

1 In the Explorer tree of the Orchestration Console, right-click *Users* > click *New User* to display the Create a New User dialog box.

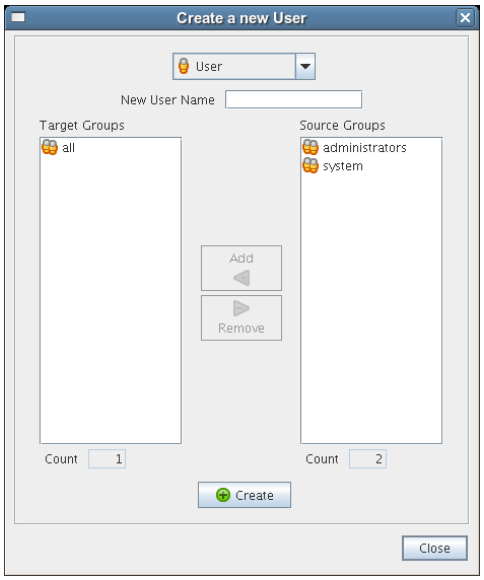


2 Specify the name of the new user you want to create in the *New User Name* field, then click *OK*.

The user account is created, but is not currently running jobs, as indicated by its object icon in the Explorer tree or in the admin view of each User group to which it belongs.

To create a new user through the *Actions* menu:

1 In the Orchestration Console, click *Actions* > *Create User* to display an expanded version of the Create a New User dialog box.



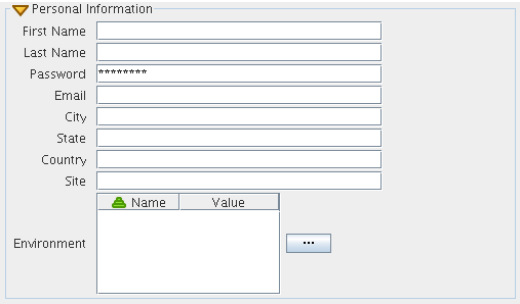
This dialog box includes a method for designating the user as a member of the *administrators* user group. In this walkthrough, we will create the user as a member of the *all* group, which does not place the user in the administrators group.

2 Specify the new username in the *New User Name* field, click *Create*, then click *Close*.

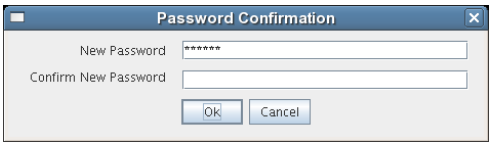
3 Define the user password.

3a In the Explorer tree of the Orchestration Console, select the new User in the Users object *all* group to open its Info/Groups page.

3b In the Info/Groups page, select the collapse/expand icon in the Personal Information section to open the fields of that section.



3c In the *Password* field, change the default password, then click the Save icon to display the Password Confirmation dialog box.



3d In the *Confirm New Password* field, enter the password you defined previously, click *OK*, then click the Save icon to save the password.

When a user logs out, the User object icon is dimmed in the Explorer tree or in the admin view of each User group to which it belongs.

1.3.4 Logging In a User for Manual Registration

If you do not select the *Auto Register Users* check box on the grid object's *Info/Configuration* page, you have the option of explicitly accepting or denying the login attempts of a user, thus preventing that user from creating an account.

1 Make sure that the *Auto Register Users* check box on the grid object's Authentication page is not selected (see Step 2 in "Automatically Registering a Resource" in the *NetIQ Cloud Manager 2.3 Installation Guide*) and that you have created a new user.

2 Use the `zos` command line interface to log in to the server.

2a From a system terminal or from an Orchestration Server login in Windows, enter the following command:

```
zos login --user=user
```

If you are attempting to log in to a machine other than the local host, you can alter the command to the following:

```
zos login Orchestration_Server_name --user=user
```

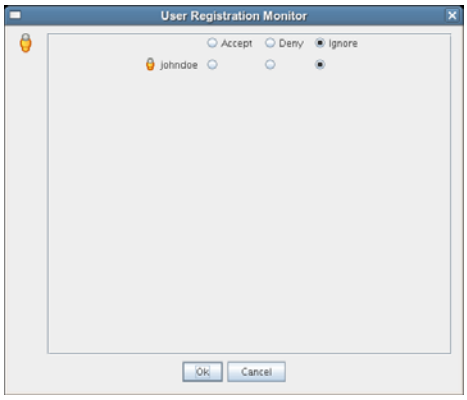
2b Enter the password for the user credentials. For this walkthrough, you can simply press *Enter* to enter an empty password.

2c When prompted for a decision regarding whether you want to accept the server certificate, enter `yes`.

An eror message is generated:

```
ERROR: login failed: user name or password 'incorrect'
```

3 In the Users Monitor, click the User Registration icon to open the User Registration Monitor dialog box.



This dialog box lets you preview the users who are trying to log in to the server. The top row of radio buttons is a mass selector for all listed users, allowing you the choice to accept, deny, or ignore automatic registration for all listed agents.

If you want to choose the users that can be allowed to auto register, you can identify the user by name and select how you want to handle that agent's request for registration the next time it tries to log in.

4 For this example, select the *Accept* radio button adjacent to the user you want to register, then click *OK*.

The user account is created, but is not currently running jobs, as indicated by its object icon in the Explorer panel, or in the Information view of each User group to which it belongs.

1.3.5 Directory Service Authentication (Optional)

There are some configuration steps you need to follow in the Orchestration Console if you want to immediately configure the authentication of both users and resources to the Orchestration Server using a directory service like ADS or LDAP. For more information, see "Authentication Page" in the *NetIQ Cloud Manager 2.3 Component Reference*.

1.4 Running a Sample Job

Deployed jobs can be run from the Scheduler utility inside of the Orchestration Server Console or from the `zos` command line. For the purpose of this walkthrough, we will run a sample job from the command line after logging in as a user.

1 From a system terminal, enter the following command:

```
zos run whoami
```

If you have more than one resource connected, you might want to run the job on one resource in particular. If you want do run the job this way, you can do so by adding an argument to the command line:

```
zos run whoami resource=name_of_resource
```

Now that you have run the sample job, you need to use some of the Orchestration Server tools to verify that it has run. For more information, see Section 1.5, “Looking at the Job After It Has Run,” on page 24

1.5 Looking at the Job After It Has Run

After you have run the job, there are several ways to verify that it has run. This section explains those methods.

- Section 1.5.1, “Verification at the Command Line,” on page 24
- Section 1.5.2, “Verification at the Jobs Monitor,” on page 25

1.5.1 Verification at the Command Line

The following sections explain some of the `zos` commands that you can use to verify that a job has run and monitor some of the results of the job:

- “zos jobs” on page 24
- “zos jobinfo job_name” on page 24
- “zos status --detail” on page 25
- “zos log job_id --verbose” on page 25

zos jobs

You can use the `zos jobs` command to list all of the jobs that have run while you have been logged in as a user. Running this command yields an output like this:

```
Job        JobID           User       Started             Elapsed  State
whoami      userA.whoami.60  userA       12/24/2008 02:35:38  0:00:00   Completed
```

zos jobinfo job_name

You can use the `zos jobinfo -e job_name` command to display information for a named job the last time it was run. Running this command yields an output like this:

```
Jobname/Parameters    Attributes
------------------    ----------
whoami              Desc: This is a demo example of enhanced exec

    numJoblets      Desc: The number of joblets to schedule
                    Type: Integer
                 Default: 1

    resource        Desc: The resource id to run on
                    Type: String
                 Default: .*
```

zos status --detail

You can use the `zos status --detail` command to display the status of the most recently run job. Running this command yields an output like this:

```
Job Status for userA.whoami.60
------------------------------
              State: Completed
     Resource Count: 0              (0 this job)
   Percent Complete: 100%
          Queue Pos: n/a
    Child Job Count: 0              (0 this job)
      Joblet Counts: 1 (0)          (1 (0/0/1/0/0) this job)

      Instance Name: whoami
           Job Type: whoami
               Memo:
           Priority: medium
          Arguments: resource=tszen4_agent

        Submit Time: 12/24/2008 02:35:38
      Delayed Start: n/a
         Start Time: 12/24/2008 02:35:38
           End Time: 12/24/2008 02:35:39
       Elapsed Time: 0:00:00
         Queue Time: 0:00:00
         Pause Time: 0:00:00

     Total CPU Time: 0:00:00        (0:00:00 this job)
      Total GCycles: 0:00:00        (0:00:00 this job)
         Total Cost: $0.0002        ($0.0002 this job)
          Burn Rate: $0.8982/hr     ($0.8982/hr this job)

   Termination Type: n/a
          Job Error: <none>

 Joblet Error Count: 0              (0 this job)
   Node Error Count: 0              (0 this job)
     Excluded Nodes: 0              (0 this job)

Bad Provision Count: 0              (0 this job)
Excluded Provisions: 0              (0 this job)
```

zos log job_id --verbose

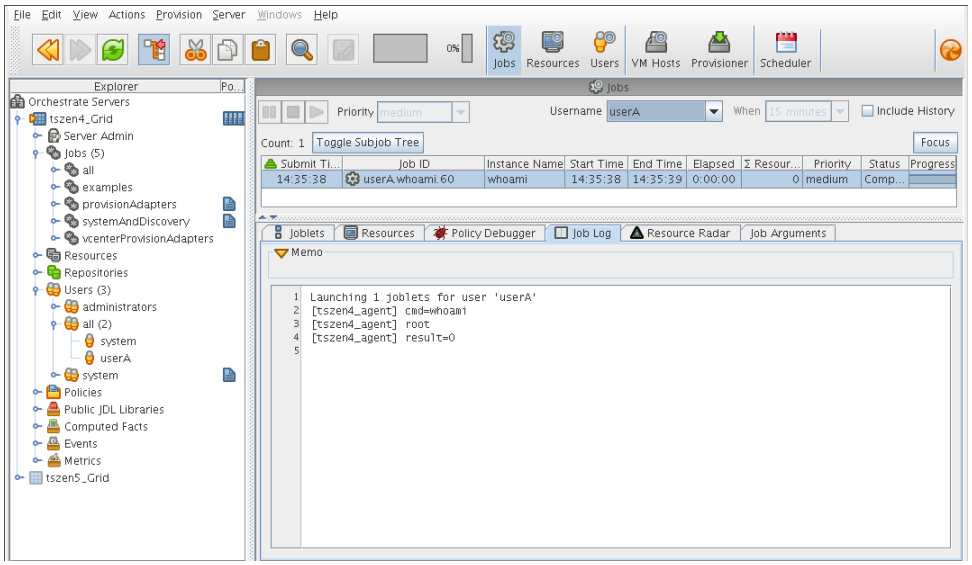
If you know the Job ID for a particular job that has run, you can use the `zos log job_id --verbose` command to display its detailed job log. Running this command yields an output like this:

```
Agent connected with ID: user_userA_64 (comms ok), Session ID: 52
Launching 1 joblets for user 'userA'
[tszen4_agent] cmd=whoami
[tszen4_agent] root
[tszen4_agent] result=0
Agent ID: null logged out.
```

1.5.2 Verification at the Jobs Monitor

If you want to use the Orchestration Server Console to verify that the job has run, you can open the *Jobs* Monitor to look at the Job Log and see the result of the job being run.

1. In the Orchestration Console, click *Jobs* in the main toolbar to open the Jobs Monitor view.
2. In the Jobs Monitor view, click the *Job Log* tab to open the Jobs page in the workspace panel.



3 In the *Username* drop-down list box, select the user who ran the `whoami` job.

Although there is much more you can learn about a job in the Jobs Monitor view, you can see by displaying a recent job that the Orchestration Console picks up the job activity and makes it available at the console.

1.6 Using the zosadmin Command to Gather Information

You can use the `zosadmin` command line to learn what users or nodes are defined in your Orchestration Server. Follow these steps to learn about the users and nodes in your system.

1 Log in to the Orchestration Server by using the following command:

```
zosadmin login
```

You can also use the server's host name as an argument when you log in.

If you use the grid= parameter, you can specify the grid name you want to log into. If other Orchestration Servers are installed on the local host, the system cannot log in to any of them unless you use this parameter. For more information, see "The zosadmin Command Line Tool" in the *NetIQ Cloud Manager 2.3 Component Reference*.

2 Enter the administrator's user name and password.

If the login was successful, the command line tool returns a message like this:

```
Logged into tszen4_grid on server 'tszen4.provo.company.com'
```

3 Enter the following command to list the active users on your Orchestration Server:

```
zosadmin users
```

You can add the `--help` option at the command line to determine the run options for this command.

4 Enter the following command to list the active nodes on your Orchestration Server:

```
zosadmin nodes
```

You can add the `--help` option at the command line to determine the run options for this command.

With the completion of this part of the walkthrough, you have a good understanding of the parts of Orchestration Server and how they work.

1.7 Stopping and Starting Orchestration Components

Use the following methods for stopping and starting Cloud Manager Orchestration components.

- Section 1.7.1, “Stopping and Starting the Orchestration Server,” on page 27
- Section 1.7.2, “Stopping and Starting the Orchestration Agent,” on page 28
- Section 1.7.3, “Starting and Stopping the Orchestration Server Console,” on page 28

1.7.1 Stopping and Starting the Orchestration Server

You can use the following methods for stopping and starting the Orchestration Server.

- “Stopping the Server” on page 27
- “Starting the Server” on page 27

Stopping the Server

You need to shut down the Orchestration Server before you power off the computer where it is running. This routine prevents possible data corruption.You must be logged in to an Orchestration Server in order to stop it. There are two methods to stop the current server:

- If you are in the Orchestration Console, click *Server* > click *Shutdown ZOS*.
- If you are at the command line, enter the following command:

  ```
  /etc/init.d/novell-zosserver stop
  ```

Starting the Server

The Cloud Manager Orchestration components installation and configuration automatically starts the Orchestration Server. The Orchestration Server must be stopped before you can start it. You must be logged in to an Orchestration Server to start it. There are two methods you can use to start the current server:

- If you are at the command line, enter the following command:

  ```
  /etc/init.d/novell-zosserver start
  ```

- To restart the Orchestration Server from the command line, enter the following command:

  ```
  /etc/init.d/novell-zosserver restart
  ```

  This command stops the server before restarting it.

1.7.2 Stopping and Starting the Orchestration Agent

You can use the following methods for stopping and starting the Cloud Manager Orchestration Agent.

- "Stopping the Agent" on page 28
- "Starting the Agent" on page 28

Stopping the Agent

There are several methods you can use to stop the agent:

- If you are in the Orchestration Server Console, select the *Resources* Monitor, select a resource, then click the red icon in the work area to shut down that agent.
- If you are running the agent on a Windows machine, click *Start* > *Programs* > *Novell* > *ZOS* > *Agent* > *Shutdown ZOS Agent*.
- If you are at the Linux bash prompt, enter the following command:

  ```
  /etc/init.d/novell-zosagent stop
  ```

Starting the Agent

- To start the Orchestration Agent from a Windows machine, double-click the *ZOS Agent* shortcut on your desktop or click *Start* > *Programs* > *Novell* > *ZOS* > *Agent* > *Start ZOS Agent*.
- If you are at the Linux bash prompt, enter the following command:

  ```
  /etc/init.d/novell-zosagent start
  ```

1.7.3 Starting and Stopping the Orchestration Server Console

You can use the following methods for stopping and starting the Cloud Manager Orchestration Console.

- "Stopping the Orchestration Console" on page 28
- "Starting the Orchestration Console" on page 28

Stopping the Orchestration Console

To stop the Orchestration Server Console at the console itself, click *File* > *Exit*, or click the shutdown icon on the console window.

Starting the Orchestration Console

To start the Orchestration Console from a Linux machine: enter `./zoc`.

- On a SLES 11 machine, if you have installed the Orchestration Server along with the Orchestration Console, change to `/opt/novell/zenworks/zos/server/bin` and enter the following command

  ```
  ./zoc
  ```

- On a SLES 11 machine, if you have installed the Orchestration Console alone, change to `/opt/novell/zenworks/zos/clients/bin` and enter the following command

  ```
  ./zoc
  ```

To start the Orchestration Console on a Windows machine, double-click the *ZOS Clients* shortcut on your desktop or click *Start* > *Programs* > *Novell* > *ZOS* > *Clients* > *Orchestration Console.*

1.8 Configuring the Orchestration Server for LDAP or ADS Authentication

There are some configuration steps you need to follow in the Orchestration Console if you want to immediately configure the authentication of both users and resources to the Orchestration Server through a directory service like ADS or LDAP.

The Cloud Manager Orchestration Server uses only one attribute of a given LDAP user: its group membership. For example, if the following settings were already configured in the Orchestration Server,

```
BaseDN 'dc=domain,dc=novell,dc=com'
UserAttribute 'uid'
UserPrefix 'ou=Users'
```

you could further configure the Orchestration Server to identify users belonging to an LDAP group using the setting `LDAP:groupnocase:administrators`.

You would do this by specifying a filter in the Orchestration Server using these settings:

```
GroupFilter 'memberUid=${USER_NAME}'
GroupPrefix 'ou=Groups'
GroupAttribute 'cn'
```

Applying these settings would let authenticated users belonging to the “administrators” LDAP group be added to the “administrators” user group in the Orchestration Server (and so allow them to log in to the Orchestration Console, for example).

For information on configuring these settings in the Orchestration Server, see “Orchestration Server Authentication Page” in the *NetIQ Cloud Manager 2.3 Component Reference*.

NOTE: Depending upon your selection at the *Server Type* drop down list on the *Enable LDAP* subpanel of the *Authentication* page of the Orchestration Console the configuration fields change to reflect the relevant settings. (One server type is *Active Directory Service*, the other is *Generic LDAP Directory Service*.)

1.9 Viewing Orchestration Server Data in the Sentinel Event Source Server

The Sentinel Event Source Server displays data from the Orchestration Server.

For each logging event, the Orchestration Server passes a free form text log message and log event metadata to the Sentinel Event Source Server. The following table shows how the data are mapped:

Table 1-1 *Orchestration Server Logged Data Mapped to Sentinel Event Source Server Fields*

Sentinel Event Field Name	Orchestration Server Data Represented
Message	A free-form text log message.

Sentinel Event Field Name	Orchestration Server Data Represented
TargetServiceName	The Orchestration Server grid name.
ReporterIP	The IP address of the Orchestration Server.
TargetServiceComp	The originating facility of the Orchestration Server event.
EventName	A name formatted as `gridname: taxonomy_key`. For more information, see "Event Classification and Taxonomy Keys" in the "*NetIQ Cloud Manager 2.3 Installation Guide*".
ProductName	The product is always `Cloud Manager Orchestration` for convenient event filtering.
InitUserName *InitUserID*	The name of the Orchestration Server user. Used for user-oriented events.

For some events, the Orchestration Server Sentinel Collector passes some structured data related to the log event. When available, this data appears in the *ExtendedInformation* Sentinel Event field, which is a list of key/value pairs delimited by semicolons. Some keys are always present in *ExtendedInformation*, but others appear only when relevant. The following table lists these keys:

Table 1-2 *Orchestration Server Key/Value Pairs Displayed in the Sentinel Event Source Server*

Key	Displayed	Value
`loglevel`	always	The original Orchestration Server log level, for example, `NOTICE`, or `INFO`.
`eventclass`	always	A symbolic identifier for this class of event, e.g. `admin_login`, `policy_association`. This also serves as the taxonomy key. The value is unclassified if the event does not contain any additional structured data.
`group`	when relevant	The Grid object group related to this log event.
`job`	when relevant	The job related to this log event.
`jobinstance`	when relevant	The job instance ID related to this log event.
`member`	when relevant	The Grid object group member related to this log event.
`policy`	when relevant	The policy related to this log event.
`repository`	when relevant	The repository related to this log event.
`resource`	when relevant	The resource related to this log event.
`schedule`	when relevant	The schedule related to this log event.
`trigger`	when relevant	The trigger related to this log event.
`type`	when relevant	The Grid object type related to this log event.
`user`	when relevant	The name of the user related to this log event. This key also appears in the *InitUserName* and *InitUserID* fields.
`vbridge`	when relevant	The Vbridge related to this log event.
`vdisk`	when relevant	The Vdisk related to this log event.

Key	Displayed	Value
vm	when relevant	The VM related to this log event.
vmhost	when relevant	The VM host server related to this log event.
vnic	when relevant	The VNIC related to this log event.
target	when relevant	The name of the Grid object related to this log event.
action	when relevant	The action related to this log event.

1.10 Orchestration Server Log Levels Mapped to Sentinel Log Levels

The information in the following table shows the correlation between the Orchestration Server logging levels and those you are likely to see in Novell Sentinel.

Table 1-3 *Orchestration Server Log Levels Mapped to Sentinel Log Levels*

Orchestration Server Log Level (alpha)	Sentinel Collector Log Level (numeric)	Comments	Usage Examples
EMERGENCY	5	Urgent conditions that require immediate attention. Indicates that the system is no longer functioning.	
ALERT	5	Conditions that should be corrected immediately.	
CRITICAL	5	Critical conditions.	
ERROR	4	Errors that have been correctly handled.	Unsuccessful job runs and provisioner/VM actions that have been correctly handled but might require manual attention.
WARNING	3	Warning messages.	Unexpected but recoverable events indicating degraded operational status, fsuch as grid objects becoming unhealthy.
NOTICE	2	Conditions that are not error conditions, but should possibly be handled specially.	Events expected occasionally as part of usual business operations, such as password changes, failed authentications and authorizations, grid objects returning to good health, server shutdown.
STATUS	1	Conditions that report on changes in operational conditions.	Events expected frequently as part of usual business operations, such as successful authentications (logins and logouts), deletion of grid objects, changes in group membership, deployment of jobs/policies/schedules/triggers, policy association.

Orchestration Server Log Level (alpha)	Sentinel Collector Log Level (numeric)	Comments	Usage Examples
INFO	1	Informational messages.	Non-security-sensitive events expected very frequently as part of usual business operations, such as successful job runs and provisioner/VM actions, resources going online/offline, session expiration/timeout.

2 Determining the Version of an Orchestration Component

Cloud Manager Orchestration is made up of many components, which you choose whether to install when you first install the product. If you are new to Cloud Manager Orchestration, if you need to determine component compatibility, or if you need to confer with NetIQ Support, it is useful to know how to obtain the version numbers of the different components.

One way to determine which version of the Cloud Manager RPM packages you have installed is to use the following command on the machine where a Cloud Manager Orchestration component is installed:

```
rpm -qa | grep novell
```

The table below provides additional methods you can use to determine the version number for Orchestration components.

Table 2-1 *Determining the Version Number for Cloud Manager Orchestration Components*

Component	How to Determine Version Number
Orchestration Server	At the Orchestration Console Explorer view, select the grid object, then open the Info/Configuration page of the workspace panel. The version number is listed in the *Server Version* field. Advanced Orchestration Server users can also find the version value in the `matrix.version` fact.
Orchestration Agent	After the agent is registered at the Orchestration Console Explorer view, select its resource object listed under *Resources*, then open the Info/Groups page of the workspace panel. On the Info/Groups page, select *Agent Information*. The version number is listed in the *Agent Version* field. Advanced Orchestration Server users can also find the version value in the `resources.agent.version` fact.
Orchestration Console	At the console, click *Help > About Cloud Manager Orchestration Console*. The console version number and license expiration date is listed on the About Cloud Manager Orchestration Console dialog box.
Command Line Tools (zos, zosadmin)	No method is currently available.

Cloud Manager Monitoring Server	On the command line of the machine where the Monitoring Server is running, enter (at `/opt/novell/zenworks/monitor/sbin`) the following command:
	
	`gmetad -V`
	
	`gmetad -- version`

Component	How to Determine Version Number

Cloud Manager Monitoring Agent	On the command line of the machine where the agent is running, enter (at `/opt/novell/zenworks/monitor/sbin`) the following command:
	`gmond -V`
	`gmond -- version`

3 Changing Orchestration Server Default Parameters and Values

The following table provides the current default values for some key performance parameters of the Cloud Manager Orchestration Server. Although the server is fine-tuned by default for optimal performance at normal loads, if you want to perform hundreds of provisioning actions simultaneously you can change some of the default settings for increased server performance in such a scenario.

Table 3-1 *Default Parameters of the Cloud Manager Orchestration Server*

Parameter Name	Shipping Default Value	Changing or Displaying the Parameter Configuration
File Descriptors Limit	2048	(Optional) You can change the value as shown below: `/etc/init.d/novell-zos-server: ulimit -n <new_limit>`
Java Heap Space	2048 MB	Create a file of the following name: `/opt/novell/zenworks/zos/server/conf/cmosadmin.properties` Add the following line to the newly created file: `default.jvmargs=-Xmx2560m`
PermGen Space	512 MB	Create a file of the following name: `/opt/novell/zenworks/zos/server/conf/cmosadmin.properties` Add the following line to the newly created file: `default.jvmargs=-XX:MaxPermSize=1536m`
Audit Queue Size Max	200	Increase the value of this parameter by using the following command: `zosadmin set --mbean="local:facility=audit" --attr=QueueSizeMax --type=Integer --value=1000`
MaxRunJobWaitTimeout	120000	You can change the value of this parameter as shown below: `zosadmin set --mbean="local:facility=broker" --attr=MaxRunJobWaitTimeout --type=Integer --value=<time_in_milliseconds>`

Parameter Name	Shipping Default Value	Changing or Displaying the Parameter Configuration
MatchingResourcesCheckinterval	30000	Increase the value of this parameter by using the following command: `zosadmin set --mbean="local:facility=broker" --attr=MatchingResourcesCheckInterval --type=Integer --value=600000`

Kernel ARP Threshold Values	◆ thresh1 = 128
	◆ thresh2 = 512
	◆ thresh3 = 1024
		Set these values higher than the default. For example:
		`cat /proc/sys/net/ipv4/neigh/default/gc_thresh1 = 256`
		`cat /proc/sys/net/ipv4/neigh/default/gc_thresh2 = 1024`
		`cat /proc/sys/net/ipv4/neigh/default/gc_thresh3 = 2048`

Job Limits: ◆ Soft Top Level Job Limit ◆ Max Queued Jobs ◆ Absolut Max Active Jobs	`matrix.maxtopjobs` = 200 `matrix.maxqueued` = 300 `matrix.maxactive` = 400	Change the Grid Object default values in the Orchestration Console as follows: `matrix.maxtopjobs` = 600 `matrix.maxqueued` = 700 `matrix.maxactive` = 1000

4 Changing Orchestration Agent Default Parameters and Values

If you notice a high incidence of joblet failures on a vsphere discovery job, the cause might be a result of a lack of available resources on the Orchestration Agent. You need to increase the Orchestration Agent heap size beyond the default, which is 512 MB.

To address this issue for an Orchestration Agent running on a Windows computer, create the `C:\Program Files (x86)\Novell\ZOS\Agent\bin\zosagent.vmoptions` file with the following content:

```
-XX:+HeapDumpOnOutOfMemoryError
-Dfile.encoding=UTF-8
-Xmx512m
-Xms512m
-XX:MaxPermSize=128m
```

This same content also exists in the `/opt/novell/zenworks/zos/agent/bin/node` file for an Orchestrate Agent running Linux.

You can adjust the values in these files as needed:

- The `-XMx` parameter sets the maximum heap size.
- The `-XMs` parameter sets the initial and minimum heap size.
- The `-XX:MaxPermSize` parameter sets the PermGen memory size.

After you create the `zosagent.vmoptions` file for the Windows agent, make configuration changes and save it, you need to restart the Windows agent.

For a Linux agent, look for the `JVMARGS=` section of the `/opt/novell/zenworks/zos/agent/bin/node` file. Normally, you change values of the `-XMx` or the `-XX:MaxPermSize` parameters in the file. After you make such changes, you need to save the file and restart the Linux agent.

NOTE: If the agent is running low on PermGen memory, we suggest that you conservatively increase it by 128 MB increments until no errors occur.

---

5 Caching Computed Facts in a Grid with Large Numbers of Resources

If your Orchestration grid includes a large number of resources with associated Computed Facts, it is likely that these computed facts are evaluated with each Ready for Work message received by the broker from the Orchestration Agent. These evaluations can cause an excessive load on the Orchestration Server, causing a decrease in performance. You might see warnings in the server log similar to the following:

```
07.07 18:27:54: Broker,WARNING: ----- Long scheduler cycle time detected -----
07.07 18:27:54: Broker,WARNING: Total:3204ms, JDL thrds:8, TooLong:false
07.07 18:27:54: Broker,WARNING: Allocate:0ms [P1:0,P2:0,P3:0], Big:488
07.07 18:27:54: Broker,WARNING: Provision:4ms [P1:0,P2:0,P3:0], Big:253
07.07 18:27:54: Broker,WARNING: Msgs:3204ms [50 msg, max: 3056ms (3:RFW)]
07.07 18:27:54: Broker,WARNING: Workflow:[Timeout:0ms, Stop:0ms]
07.07 18:27:54: Broker,WARNING: Line:0ms, Preemption:0ms, (Big: 3), Mem:0ms
07.07 18:27:54: Broker,WARNING: Jobs:15/0/16, Contracts:10, AvailNodes:628
07.07 18:27:54: Broker,WARNING: PermGen: Usage [214Mb] Max [2048Mb] Peak
[543Mb]
07.07 18:27:54: Broker,WARNING: Memory: free [1555Mb]  max [3640Mb]
07.07 18:27:54: Broker,WARNING: Msgs:483/50000 (recv:128692,sent:14202),
More:true
07.07 18:27:54: Broker,WARNING: ------------------------------------------------
```

To work around this issue, we recommend that you cache the Computed Facts.

1. In the Explorer tree of the Orchestration Console, expand the Computed Facts object, then select *vmbuilderPXEBoot*.

   The `vmbuilderPXEBoot` fact does not change, so setting the cache here is safe from any automatic modifications.
2. In the Computed Facts admin view, select the *Attributes* tab to open the Attributes page.
3. In the Attributes page, select the *Cache Result for* check box, then in the newly active field, enter 10 minutes (remember to change the drop-down list to indicate *Minutes*).

   This value must be greater than the default of 30 seconds.
4. Click the *Save* icon to save the new configuration.

NOTE: If necessary, you can also cache other computed facts to improve server performance.

5.1 Walkthrough: Scheduling a System Job

This section demonstrates how you can use the Orchestration Console to deploy and schedule an existing system job named `auditCleaner.job`. This example job is included in the `examples` directory of your Orchestration Server installation.

This section includes the following information:

- Section 5.1.1, “Deploying a Sample System Job,” on page 40
- Section 5.1.2, “Creating a New Schedule for the Job,” on page 41
- Section 5.1.3, “Defining the New Schedule,” on page 42

- Section 5.1.4, “Activating the New Schedule,” on page 46
- Section 5.1.5, “Running the New Schedule Immediately,” on page 46

5.1.1 Deploying a Sample System Job

Before a job can run, the Orchestration Server administrator must deploy it, which involves moving it from a developed package state to a state where it is ready and available for users. Only the administrator has the necessary rights to deploy a job.

There are four methods you can use to deploy a job:

- Deploy it from the Orchestration Console by right-clicking the *Jobs* container in the Explorer panel.
- Deploy it from the Orchestration Console by selecting the *Actions* menu in the console.
- Deploy it from the `zosadmin` command line (`zosadmin deploy path_to_job`).
- Copy the deployable component to the “hot” deployment directory under the Orchestration Server instance directory. Typically, this directory is located at `/var/opt/novell/zenworks/zos/server/deploy`. Using this method, deployment proceeds within a few seconds. The server monitors this directory.

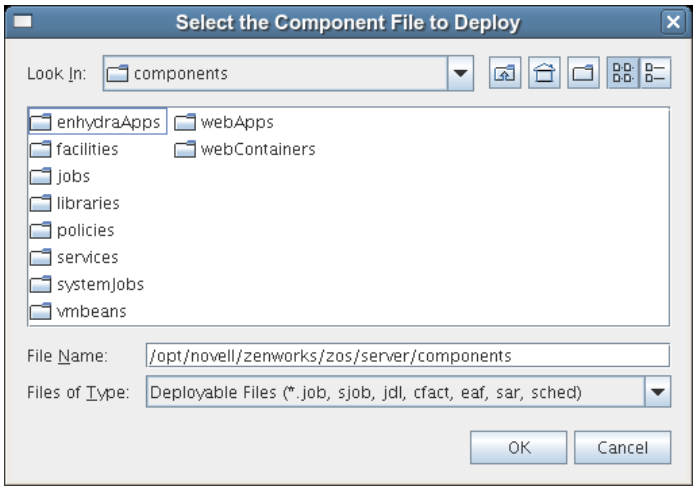
A runnable job can also be scheduled, which means that the schedule for running the job and the trigger or triggers that initiate or “fire” the schedule (or both) are configured and packaged with the job.

For this walkthrough, you deploy one of several system jobs (`auditCleaner.job`) developed for Cloud Manager Orchestration Server administrators to demonstrate how system jobs are deployed and run. This job package, which is actually a `.jar` archive, includes only a `.policy` component and a `.jdl` component. It does not have a `.sched` component. You can use the Job Scheduler in the Orchestration Console to add the `.sched` component separately.

NOTE: A job developer can create and package jobs that include a `.jdl` file, a `.policy` file, a `.trig` file (trigger), and a `.sched` file (schedule). The presence of the `.sched` file in the job package is also typical of the predeployed discovery jobs installed with the Orchestration Server, which run without intervention when the criteria for firing the schedule are satisfied. Such jobs are visible in the Job Scheduler because they already include the `.sched` components.

Although this walkthrough demonstrates only the first method listed above for deploying, the other methods are relatively simple, so no further examples are provided to illustrate them.

1. In the Explorer panel of the Orchestration Console, right-click the *Jobs* container, then click *Deploy Job* to open the Select the Component File to Deploy dialog box.

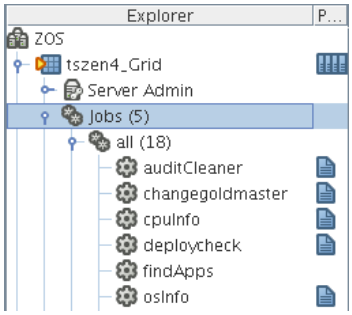


2 Open the *Look In* drop-down list, then navigate to the location of the job you want to deploy.

Although a job developer can store jobs at any location on the network, the sample jobs shipped with the Orchestration Server are limited to the directories where the product is installed. For this walkthrough, navigate to the `/opt/novell/zenworks/zos/server/components/systemJobs` directory.

3 Select *auditCleaner.job*, then click *OK* to deploy the job to the *Jobs* container.

The job appears in the *all* container and in the *examples* container in the tree.



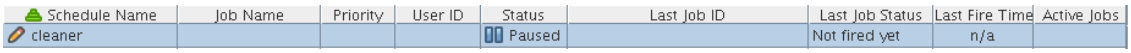
5.1.2 Creating a New Schedule for the Job

When a job has been deployed, you can create a schedule to specify when you want it to run. In this walkthrough, you create a schedule for the `auditCleaner` job by using the Scheduler tool in the Orchestration Console.

1 In the toolbar in the Orchestration Console, click the Job Scheduler button to open the Job Scheduler view.

2 In the Job Scheduler view, click *New* to open the Enter Unique Schedule Name dialog box.

3 Specify a name for the schedule you want to create for this job. For this walkthrough, specify the name `cleaner` in the *Schedule Name* field, then click *OK* to return to the Job Scheduler view.



The new schedule is highlighted in the Job Schedules Table and is flagged with a pencil icon, signifying that the schedule has not been committed yet. Continue with Section 5.1.3, “Defining the New Schedule,” on page 42 to define this new schedule by adding the specific information you want.

5.1.3 Defining the New Schedule

Defining a new job schedule consists of selecting its general properties, its specific properties, and the triggers you want to be associated with it.

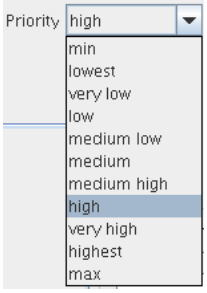
- “Choosing General Properties for a New Schedule” on page 42
- “Creating and Assigning a Time Trigger for the New Schedule” on page 43
- “Adding Specific Parameters to the New Schedule” on page 44

Choosing General Properties for a New Schedule

After you have created a new job schedule, its name cannot be changed, but you can add properties to it that help to identify and classify it in a general way. Use the following steps to add these properties:

1. In the Job Schedule Editor panel of the Scheduler view, select the *Job* drop-down list.
2. From the list of available jobs, select *auditCleaner* as the job to which this schedule applies.
3. In the Job Schedule Editor, select the *User* drop-down list.
4. From the list of available users, select *zosSystem* as the user who runs this job.

   The zos user is the built-in user that is always present. It is commonly used for routine jobs like this example.
5. In the Job Schedule Editor, select the *Priority* drop-down list.



6. From the list of available priorities, select *high* as the priority for this job schedule.

   The maximum selectable priority is dependent on an attribute associated with the selected user.
7. Click the *Save* button on the toolbar of the Orchestration Console to save the general properties you have selected for the new schedule.

   The schedule is now committed, and the attribute columns in the Job Schedules Table are populated with the name of the job that the schedule will run, the user it will run as, the priority at which it will run, and its current status. Because the schedule has not been activated yet, it remains in a *Disabled* state.

When you have chosen the general properties of the new schedule, you can either continue with “Adding Specific Parameters to the New Schedule” on page 44 or proceed directly to “Creating and Assigning a Time Trigger for the New Schedule” on page 43.

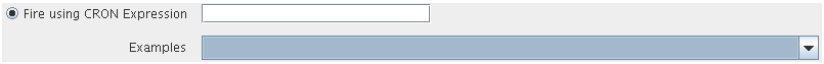
Creating and Assigning a Time Trigger for the New Schedule

A job already defined in a schedule can be triggered with two main themes: the occurrence of an event or the arrival of a point in time. In this walkthrough, you define a time trigger for the `cleaner` schedule.

In this example, there are no defined time triggers in the Job Scheduler, so you use the following steps to define a time trigger.

1. In the Job Schedule view, click *Edit Triggers* to display the Triggers dialog box.

   Time triggers are shareable across schedules. After a time trigger is defined, it is added to a list of triggers in this dialog box. You can select a predefined trigger from this list when you create a new schedule, or you can create a new time trigger, as the next steps demonstrate.
2. In the Triggers dialog box, click *New* to clear and activate the fields in the dialog box for the creation of a unique time trigger.
3. In the Enter Unique Trigger Name dialog box, specify `24 hour` as the unique name of this time trigger, then click *OK*.
4. In the *Description* field, specify `Runs every 24 hours at noon` as the description for this time trigger.
5. Click *Fire Using CRON Expression* to activate the fields for defining a cron expression.



6. Click the drop-down list of sample cron expressions, then select the default cron expression, `0 0 12 * * ?`, which is listed first.

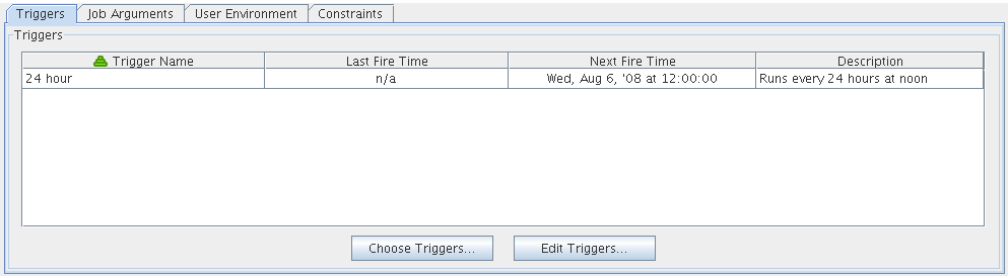
   The sample expressions in the drop-down list show cron strings with accompanying descriptions to remind you how a cron string is constructed. The examples are selectable and editable and can be used in the schedule, just as you have done in this step.

   NOTE: For detailed information about cron syntax, see “Understanding Cron Syntax in the Job Scheduler”in the *NetIQ Cloud Manager 2.3 Component Reference*.

   ---
7. Click *Save* to save the trigger you just created, then click *Close* to return to the Job Scheduler view.
8. From the Job Scheduler view, make sure that the *cleaner* schedule is selected, then click *Choose Triggers* to display the Choose Triggers dialog box.
9. In the Choose Triggers dialog box, select *24 hour* (the name of the trigger you created), click *Add* to move the trigger definition to the *Scheduled Job Triggers* list, then click *OK* to return to the Job Scheduler view.

   NOTE: You can select and combine as many time triggers as you want to apply to a given schedule. You can also combine time triggers with event triggers on a given schedule.

   In the *Triggers* list of the Job Scheduler view, the *24 hour* trigger is now associated with the new schedule.



10 Click the *Save* button to update the Orchestration Server with the new schedule/trigger association.

Adding Specific Parameters to the New Schedule

If necessary, you can now add specific parameters to the schedule to edit its job arguments, to choose whether you want to pass the user environment variables to the schedule, or to specify policy constraints to further focus the purpose of this schedule when it fires.

For the purpose of this walkthrough, none of these specific parameters is modified, although a general overview of how to do so is explained.

The following specific parameters can be managed in the Job Scheduler Editor:

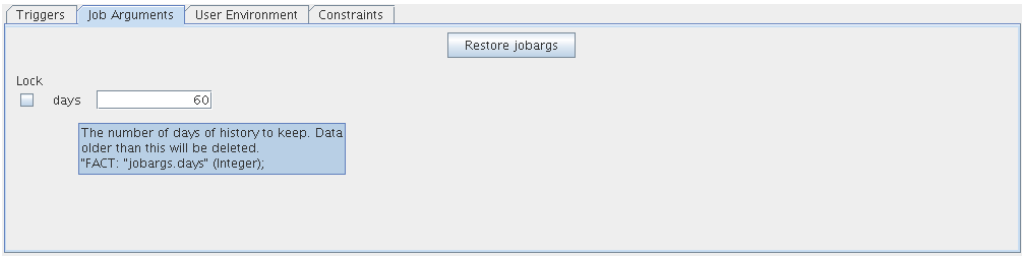
- “Job Arguments” on page 44
- “User Environment” on page 45
- “Constraints” on page 45

Job Arguments

As explained in “Creating or Modifying a Job Schedule” in the *NetIQ Cloud Manager 2.3 Component Reference*, a job argument defines the values that can be passed in to the process when a job is invoked. These values let you statically define and control job behavior. The job arguments that appear in the *Job Arguments* tab of the Schedule Editor depend on the job. The job might have no arguments.

By default, the `auditCleaner` job lists only one job argument, *jobargs.days*.

Figure 5-1 *The Job Arguments Tab of the Job Schedule Editor*



According to the tooltip text, this argument is the number of days of job history to keep, so this job cleans up the history of the job in the Orchestration Server audit database after the job reaches the age of 60 days. Data older than 60 days is to be deleted. If you want to, you can change this parameter, or any other parameter in a job argument.

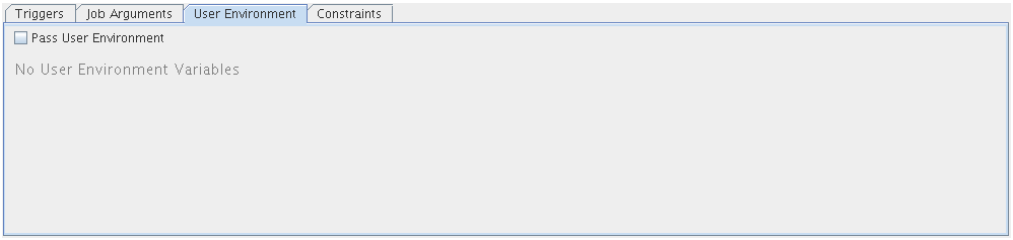
If the default value for a job argument parameter is missing, the job might fail, so you should inspect these parameters carefully.

User Environment

As explained in "Creating or Modifying a Job Schedule" in the *NetIQ Cloud Manager 2.3 Component Reference*, a user's environment variables are available in the Job Scheduler only if that user utilizes the zos command line tool and elects to pass those environment variables to the server at login time or when he or she runs a job (running the job creates the environment variables as facts in the job). The `zos run` command passes the environment for that particular job run only.

In this walkthrough, the `zosSystem` user shows no user environment variables.

Figure 5-2 *The User Environment Tab of the Job Scheduler Editor, No User Environment Variables Available*

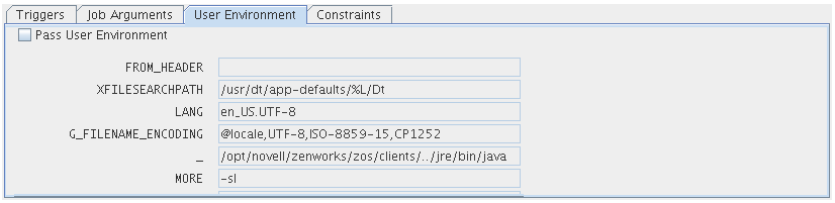


Because there are no environment variables listed, there are none to pass to the schedule, so it is not necessary to select the *Pass User Environment* check box. By default, this check box is not selected, even if environment variables are present for a user specified to run the job.

Sometimes a job is written to work from a user's environment variables. In this case, if a user has not logged in or has not run the job from the zos command line using the necessary environment option, the schedule must pass those variables to the job when it is invoked.

If you associate a user who has user environment variables with this schedule, you would see a list of those environment variables as they would be passed to the job.

Figure 5-3 *The User Environment Tab of the Job Schedule Editor, User Environment Variables Available*



Selecting the *Pass User Environment* check box in this scenario would create these variables as facts used for this job invocation.

Constraints

As explained in "Creating or Modifying a Job Schedule" in the *NetIQ Cloud Manager 2.3 Component Reference*, the *Constraints* tab displays a constraint editor that you can use to create additional constraints for the job being scheduled.

Any other constraints associated with the context of this job invocation (including but not limited to this job), with the user you've selected, with that user's group, with the jobs group, with the resources that the job uses, or with the resource groups that the job uses, run in spite of the policy that you define here. These additional constraints usually restrict or refine what the job does when this schedule fires.

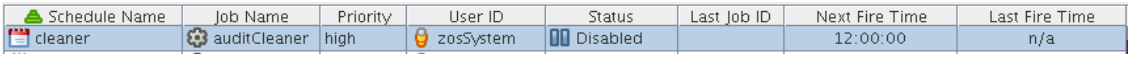
These constraints are passed to the job only when this schedule is invoked. For example, you could add a start constraint to delay the start of a job, a resource constraint to run on only one of three named machines, or a continue constraint to automatically time out the job if it takes too long to run. Anything you can do with a regular job policy constraint, you can add as a special constraint here for this particular schedule invocation.

Click the *Save* button to update the Orchestration Server with the new schedule.

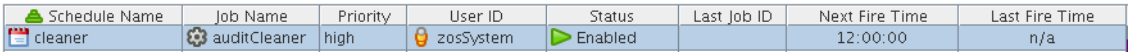
5.1.4 Activating the New Schedule

When the new schedule has been created and its triggers defined, you need to take it from the disabled state to an active state where it is ready to run.

1 In the Job Scheduler view, select the newly created job. The job shows that it is in a *Disabled* state.



2 Click *Enable* to enable the schedule.

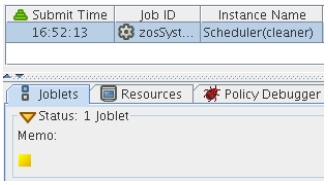


The schedule is now enabled, but has not run yet.

5.1.5 Running the New Schedule Immediately

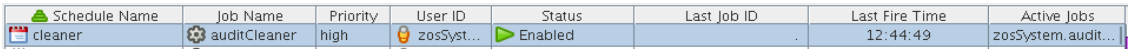
You can trigger the schedule immediately, rather than waiting for the triggers to fire.

1 In the Job Schedules Table of the Job Scheduler view, select *cleaner* (the name of the schedule you want to run), click *Run Now*, then click the job monitor button on the toolbar (*Jobs*) to open the Job Monitor view.



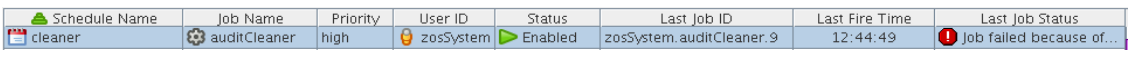
The joblet icon shows that the job is running.

2 Click the Job Scheduler button on the toolbar to open the Job Scheduler view.



The `cleaner` schedule is listed as an active job. This indicates that the schedule has started the job as anticipated.

If you click the *Refresh* button, you can see that the job now has a Job ID.



If the job invocation fails, as in this example, a red exclamation icon is also displayed.

6 Managing Virtual Machine Hosts

After you install the Cloud Manager Orchestration Agent on a VCenter Server, you need to associate the vsphere_client policy to a vSphere resource before the Discover VM hosts job can work.

- Section 6.1, “Discovering VM Hosts and Repositories,” on page 47
- Section 6.2, “Discovering VM Images,” on page 48
- Section 6.3, “Resynchronizing the VM Host’s State,” on page 49
- Section 6.4, “Shutting Down VM Hosts,” on page 49
- Section 6.5, “Restarting VM Hosts,” on page 50
- Section 6.6, “Understanding VM Host Failover,” on page 51

6.1 Discovering VM Hosts and Repositories

1. Ensure that the policies appropriate to the VM technology are configured. For more information, see “Orchestration Provisioning Adapters” in the *NetIQ Cloud Manager 2.3 Component Reference*.

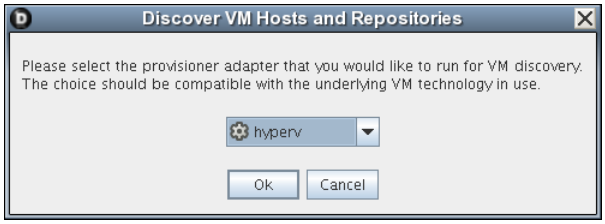
   For vSphere, the default number of slots is 4. We recommend that you increase this number to 10, depending on the hardware and available computing resources (RAM/CPU) of the server where the agent is running (this is also the server associated with the `vsphere_client` policy). Each joblet slot causes a separate Java instance on this resource. Each Java instance uses 50-60 MB of RAM and is quite CPU-intensive.

   For more information on the policies, see “Orchestration Provisioning Adapters” in the *NetIQ Cloud Manager 2.3 Component Reference*.

2. Ensure that you have set the correct number of joblet slots for the VM hosts in the policies appropriate to the VM technology. For more information, see *Joblet Slots* in the *NetIQ Cloud Manager 2.3 Component Reference*.

3. In the Orchestration Console, click *Provision* > *Discover VM Hosts and Repositories*.

   The Discover VM Hosts and Repositories dialog box is displayed.



4. Select your provisioning adapter from the drop-down list.

5. Click *OK*.

6. Click *Jobs* to view the *Jobs* section in the Orchestration Console and verify that the job has started.

   After your VM host machines are discovered, you can refresh your tree view or wait for the automatic tree refresh to see the VM host machine listed under the provisioning adapter, although no VMs are listed.

This also discovers:

- Local repositories for all types of hypervisors.
- SAN and NAS repositories for vSphere.

  To view the discovered repositories, click *Repositories*, then click *esx*.

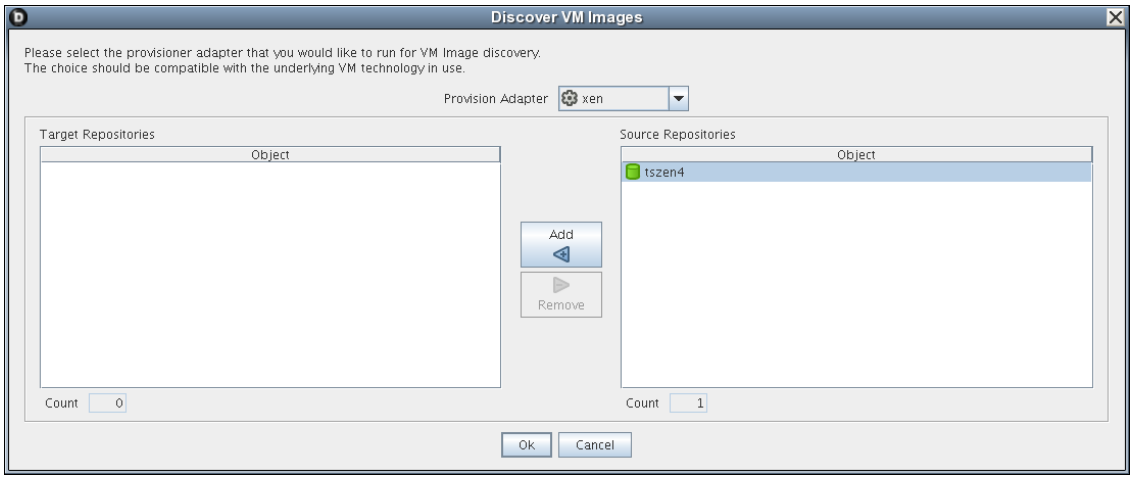
For a list of the supported guest operating systems supported by both Cloud Manager and vSphere, see "Requirements and Cloud Manager Support for the Virtual Environment" in the *NetIQ Cloud Manager 2.3 Installation Guide*.

6.2 Discovering VM Images

To discover the VM images on a specific repository:

1 In the Orchestration Console, click *Provision* > *Discover VM Images*.

  The Discover VM Images dialog box is displayed.



2 In the *Provisioning Adapter* drop-down list, select the provisioning adapter for which you want to discover the VM images.

  The source repositories for the selected provisioning adapter are displayed.

  For information on provisioning adapters, see Section 7.1, "Provisioning a Virtual Machine," on page 53.

3 Select the source repositories, then click *Add*.

  The selected repositories are added to the *Target Repositories* pane.

4 Click *OK*.

  The VM images are discovered: a separate job is launched for each repository that the user selected. After all jobs are complete, you can refresh the Explorer tree in the Orchestration Console to see the discovered VMs.

6.3 Resynchronizing the VM Host's State

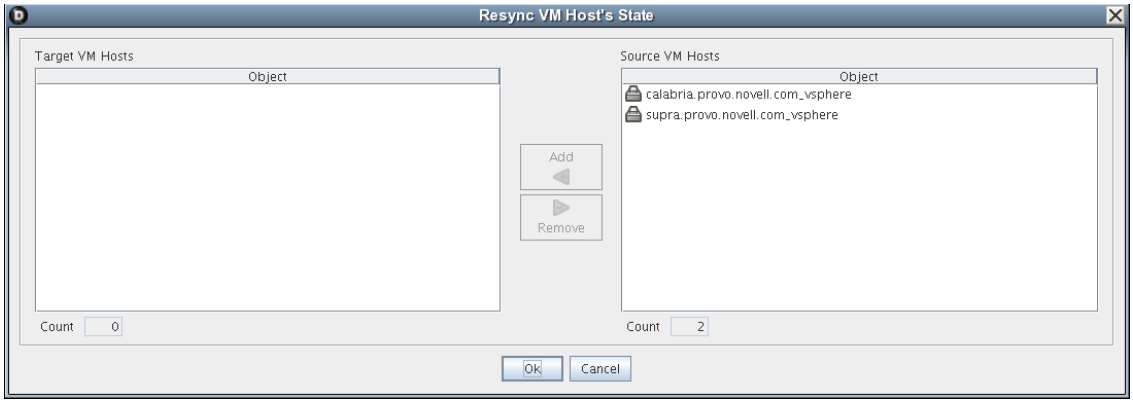
To manually verify and ensure that the state of a VM host displayed in the Orchestration Console is accurate:

1 In the Orchestration Console, right-click the VM host, then click *Discover*.

To manually verify and ensure that the state of multiple VM hosts displayed in the Orchestration Console is accurate:

1 In the Orchestration Console, click *Provision* > *Resync VM Host's State*.

The Resync VM Host's State dialog box is displayed.



2 In the *Source VM Hosts* pane, select the VM hosts to be resynchronized, then click *Add*.

The selected VM hosts are added to the *Target VM Hosts* pane.

3 Click *OK*.

6.4 Shutting Down VM Hosts

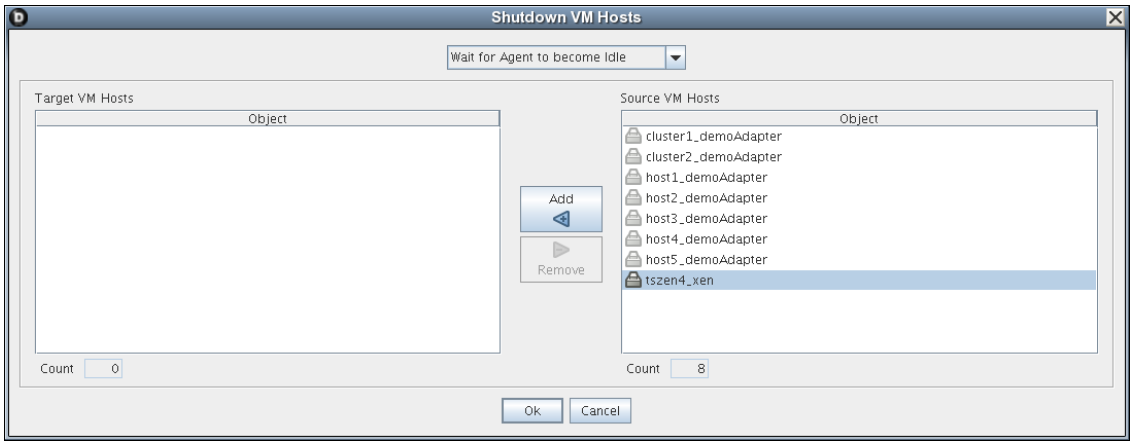
To shut down a single VM host:

1 In the Orchestration Console, right-click the VM host you want to shut down, then click *Shutdown*.

To shut down multiple VM hosts:

1 In the Orchestration Console, click *Provision* > *Shutdown Hosts*.

The Shut Down VM Hosts dialog box is displayed.



2 Choose when to shut down the VM hosts.

You can choose to shut down the VM hosts after the Orchestration Agent becomes idle or to immediately shut down the VM hosts. By default, the *Wait for Agent to become Idle option* is selected.

3 In the *Source VM Hosts* pane, select the VM hosts you want to shut down, then click *Add*.

The selected VM hosts are added to the *Target VM Hosts* pane.

4 Click *OK*.

The VMs running on the host are automatically shut down and the VM host is moved to the Shutting Down state, where will not accept any Provisioning actions.

6.5 Restarting VM Hosts

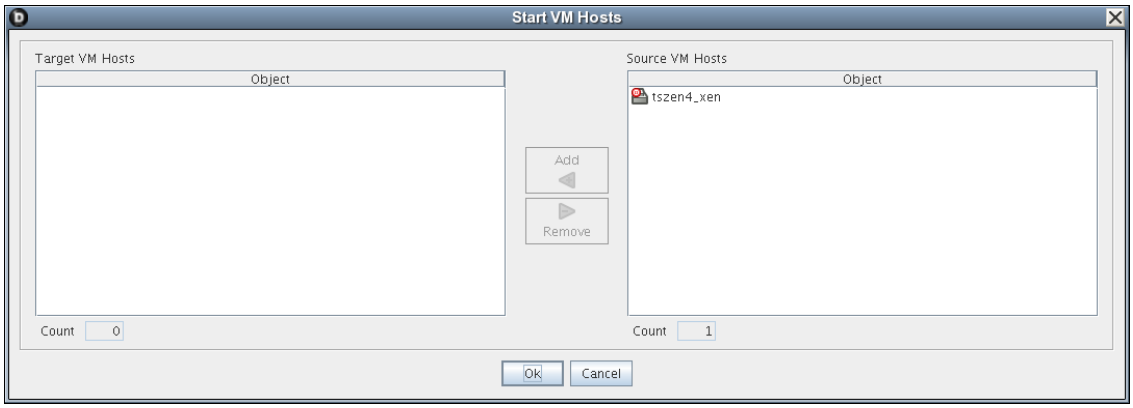
To restart a single VM host:

1 In the Orchestration Console, right-click the VM host you want to start, then click *Start*.

To restart multiple VM hosts:

1 In the Orchestration Console, click *Provision* > *Start VM Hosts*.

The Start VM Hosts dialog box is displayed.



2 In the *Source VM Hosts* pane, select the VM hosts you want to restart, then click Add.

The selected VM hosts are added to the *Target VM Hosts* pane.

3 Click *OK*.

6.6 Understanding VM Host Failover

When the Orchestration Server comes back online after being offline, it rediscovers the state of all resources, including VM hosts and the VMs running on those hosts. This section provides more information about how the Orchestration Server behaves when the VM Host loses its agent connection.

There are two possible scenarios that can occur when a VM Host fails while running VMs. The failover behavior depends on where the VM image is stored and whether the VM has the agent installed.

The following table shows possible failover scenarios with the VM Host and the expected server behavior when it occurs.

Table 6-1 *Orchestration Server Behavior when the VM Host Loses Its Agent Connection*

Scenario	Failover Behavior
Scenario 1: The VM image is: ◆ Stored on a non-local repository (for example, the zos repository) ◆ Accessible by other VM hosts ◆ Successfully provisioned **Situation:** The VM host fails.	The VMs that had been running on the failed VM host are reprovisioned to other available VM hosts. ◆ If the VM was provisioned from a template, there is now another instance of the VM. For example, if the template name is `sles10template`, the original VM provisioned from the template is then named `sles10template-1`. If the host running `sles10template-1` goes down, or if it loses its agent connection, a new instance of the template named `sles10template-2` is reprovisioned to an available host. ◆ If the original VM was a standalone VM, it is reprovisioned to an available host.
Scenario 2: The VM image is stored on a local repository. **Situation:** The VM host loses its agent connection.	◆ Because the VM image is stored locally, it cannot be reprovisioned to another VM host. ◆ When the VM host comes back online, it is reprovisioned to the host where it is stored.

In either of these scenarios, if the Orchestration Agent is installed on the VM and if the VM host loses its agent connection but the VMs retain their agent connection (for example, if someone kills the agent process on the VM host), no reprovisioning occurs.

If the VM host loses its agent connection, and if the Orchestration Agent is not installed on the running VMs, the VMs can continue running indefinitely. However, if the location of the VM image warrants it, the VMs are reprovisioned to other available hosts. When there are two (or more) of the same VM instance running on different VM hosts, the Orchestration Server is aware only of the VMs running on a VM host with an active agent connection, so the administrator must stop the VMs on the host that has lost its agent connection.

NOTE: If you are interested in failover in a high availability environment, see "Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support" in the *NetIQ Cloud Manager 2.3 Installation Guide*

7 Managing Virtual Machines

Review the following sections for specific information about the VM management functions in Cloud Manager Orchestration:

- Section 7.1, “Provisioning a Virtual Machine,” on page 53
- Section 7.2, “Provisioning Actions and History,” on page 54
- Section 7.3, “Managing a Running Virtual Machine,” on page 57
- Section 7.4, “Resynchronizing the State of All VMs,” on page 63
- Section 7.5, “Resynchronizing the State of All VMs of a Specific VM Host,” on page 63
- Section 7.6, “Shutting Down Multiple VMs,” on page 64
- Section 7.7, “Destroying and Deleting a Virtual Machine,” on page 64

7.1 Provisioning a Virtual Machine

Provisioning is the first step in a VM’s life cycle. The Orchestration Server determines the best VM host machine for running the VM, unless you select a specific host server, datastore, or network to run the VM.

By default, you can run eight VMs at one time on a VM host. If you want to provision additional VMs, you must proportionately increase the `vmhost.maxvmslots` fact value for a particular VM host in the Orchestration Console.

Provisioning VMs that have only an NPIV disk is not supported. You can provision a VM that has a hard disk and an NPIV disk (that is, a SAN repository). The OS image of the VM is stored on the local hard disk and the data resides in the SAN repository.

The Orchestration Server uses provisioning adapters to perform life cycle functions. Provisioning adapters are programs that control (start, stop, snapshot, migrate, or pause) a VM. They run as regular jobs on the Orchestration Server.

The system can discover SAN repositories for vSphere hosts.

The constraints used to determine a suitable VM host evaluate the following criteria to provide heterogeneous VM management:

- Machine architectures
- CPU
- Bit width
- Available virtual memory
- Other administrator-configured constraints, such as the number of virtual machine slots

For procedures and more information on provisioning VMs, see Section 7.3, “Managing a Running Virtual Machine,” on page 57.

7.2 Provisioning Actions and History

The following information is included in this section:

- Section 7.2.1, “What are Provisioning Actions?,” on page 54
- Section 7.2.2, “How Actions Are Displayed in the Orchestration Console,” on page 54
- Section 7.2.3, “Prioritized Provisioning Actions,” on page 57

7.2.1 What are Provisioning Actions?

The provisioning operations you perform in Cloud Manager Orchestration are recorded as “actions.” For example, in the Orchestration Console main menu:

- A VM Host Discovery action is initiated if you select *Provision* > *Discover VM Hosts and Repositories* and then you select a provisioning adapter in the Discover VM Hosts and Repositories dialog box.
- A VM Discovery action is initiated if you select *Provision* > *Discover VM Images* and then you select a provisioning adapter in the Discover VM Images dialog box. An action is specified for each Repository you specify.
- A Migrate action is initiated when you perform the migration of a Virtual Machine (VM).

For a comprehensive list of the provisioning operations supported by each Cloud Manager Orchestration provisioning adapter, see “Orchestration Provisioning Adapters” in the *NetIQ Cloud Manager 2.3 Component Reference*.

7.2.2 How Actions Are Displayed in the Orchestration Console

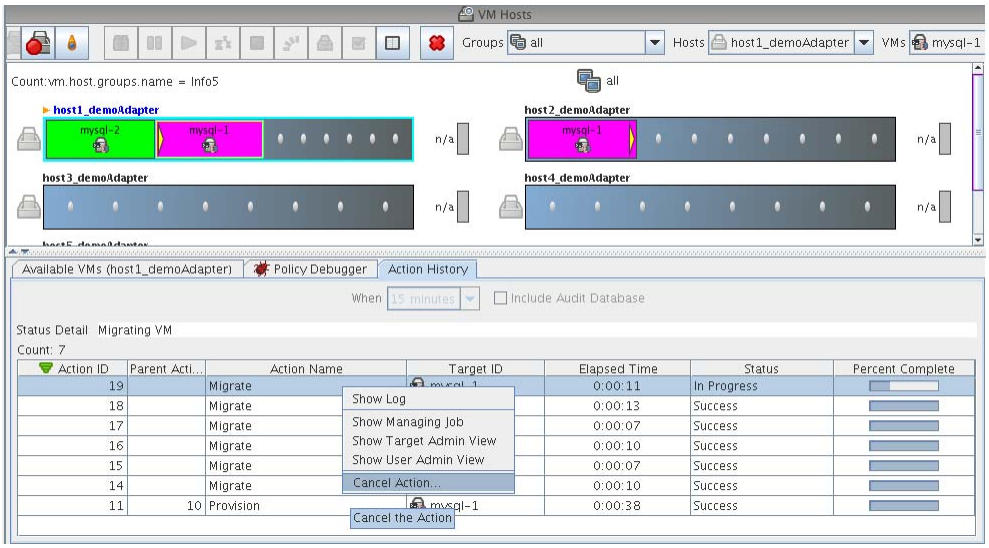
Depending on the Grid object you select, an Action History tab is displayed in several views of the Orchestration Console.

- “Action History in Monitor Views of the Orchestration Console” on page 54
- “Action History in Admin Views of the Orchestration Console” on page 55

Action History in Monitor Views of the Orchestration Console

You can see the Action History tab in the VM Hosts monitor view if you select a migrating VM:

Figure 7-1 *VM Hosts Actions*



Two action-specific menu selections are available if you right-click an action in the action history table:

- *Show Log* opens the provision log for the VM
- *Cancel Action* cancels the selected active action

The action history table is updated at the same time the polling view is updated.

NOTE: The format of action history table is similar in the Provisioner monitor view and in the Users monitor view.

If the *Include Audit Database* check box is selected in this view, the action status is not polled. Click the refresh icon to fetch and display fresh data.

NOTE: The Orchestration Server must be connected to an audit database for the *Include Audit Database* check box to be available. This behavior is the same in the Job monitor view.

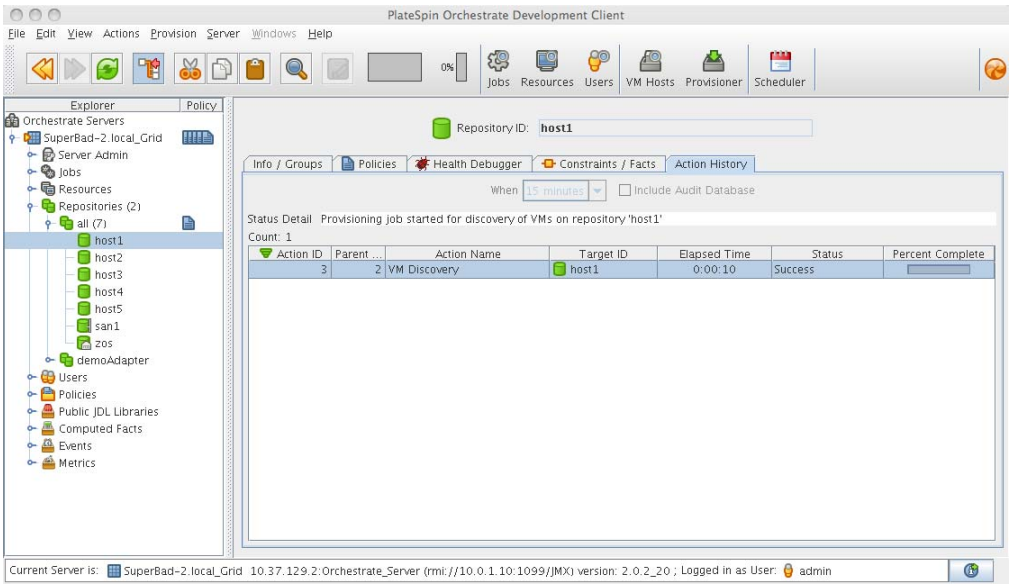
Action History in Admin Views of the Orchestration Console

Action history is displayed in other Grid object admin views of the Orchestration Console:

- User object
- Resource object
- Repository object

The following illustration shows an example of action history in the Repository admin view:

Figure 7-2 *Repository Action History*



The table below defines some of the column names and the values that can populate those columns in the action history table:

Table 7-1 *Action History Table Columns*

Column Name	Purpose
Action ID	A unique integer value used to identify the action.
Parent Action ID	If a value is displayed, the action is a child of this parent, identified with a unique integer value.
Action Name	Specific to the action being invoked.
Target ID	The identifying string of the object where the action is to occur.
Status	Displays the status of this action (specific to the action being invoked). Possible values in this column include: ◆ Started ◆ In Progress ◆ Success ◆ Failed ◆ Canceled

7.2.3 Prioritized Provisioning Actions

Because the Orchestration Server must efficiently manage many resources and jobs while providing acceptable response times, short-running, user-facing jobs, such as stopping a running VM instance are prioritized above long-running background tasks such as building a VM image.

The Orchestration Server increases the priority for the *Start*, *Stop*, *Pause*, *Suspend, CheckStatus* and *Resume* provisioning actions so that they are run before the *Clone*, *Build*, or *Save Config* provisioning actions. These prioritized provisioning actions run as “privileged provisioning actions,” which means that they can run in a resource’s extra system joblet slots. This is done so that customer-initiated provisioning actions can be run immediately, even when resource joblet slots are filled by lower priority actions, congesting the grid.

NOTE: A prioritized provisioning action runs with a new “mediumlow” priority in the Orchestration Console rather than the traditional “low” priority.

For more information about adding joblet slots, see Joblet Slots in the “The Resource Object”section of the *NetIQ Cloud Manager 2.3 Component Reference*.

7.3 Managing a Running Virtual Machine

There are many ways you can control a VM after it has been deployed. All actions from provisioning to shutting down the VM can be managed directly from the Orchestration Console, through provisioning adapter jobs, and through custom jobs written by the user.

Review the following sections for ways to manage running VMs:

- Section 7.3.1, “Using the Right-Click Menu for Provisioning Actions,” on page 57
- Section 7.3.2, “Releasing a Virtual Machine from Usage,” on page 61
- Section 7.3.3, “Managing Virtual Machine Templates,” on page 61
- Section 7.3.4, “Managed Virtual Machine Actions,” on page 62

7.3.1 Using the Right-Click Menu for Provisioning Actions

You can perform provisioning actions by right-clicking a VM in the Explorer tree of the Orchestration Console.

For information on provisioning adapters, see Section 7.1, “Provisioning a Virtual Machine,” on page 53.

The provisioning actions available from the right-click menu are as follows:

Table 7-2 *Right-Click VM Commands*

Action	Description
Provision	Starts a VM to a running state. The Orchestration Server automatically looks for the best VM host machine to run the VM on, unless you have specifically designated another server to run the VM. If a VM has snapshots, you cannot start the VM on a different host. If a VM that has snapshots is on a shared repository, you can register the VM to a different host and start the VM if the host is also connected to the shared repository.
Pause	Stops the current state of a VM and prevents it from gaining access to the processor of the host machine, although it is still resident in the memory of the host machine. Effectively, this releases the RAM and CPU resources currently being used by the VM. The current state of the VM is not saved when you pause a VM.
Resume	Allows a paused VM to access the processor of the host machine again.
Suspend	Takes a snapshot of the VM's disk and memory status, saving its current state. In the suspended state, a VM can be moved or migrated to another host machine. **NOTE:** A suspended VM must be provisioned to make it active again. The *Resume* action does not perform this function.
Shutdown	Stops a VM from running, just like shutting down a physical machine. The operating system stops and acts as if it is shut down.
Restart	Shuts down and restarts a running VM.
Migrate	**vSphere:** Migrates the VM from one host machine to another only if both the source and destination host machines have VMotion enabled. VM migrations can be of the following types: ◆ A "warm migrate" is the migration of a suspended VM to another host. From a user's perspective, if a VM is suspended, it is effectively "down." This function requires shared storage. ◆ A "hot migrate" (also called a "live migrate") is the migration of a running VM to another host and starting it there with minimal resulting downtime (measured in milliseconds). This function requires shared storage. For more information, see the VMware VMotion documentation (http://www.vmware.com/products/vmotion/).
Resync State	Ensures that the state of the VM displayed in the Orchestration Console is accurate.
Rediscover VM	Not supported on VMs managed by vSphere.
Save Config	Requests that the runtime VM configuration be persisted on the VM image for future use. **NOTE:** If you add a vNIC or a vDisk to a VM, you must initiate the *Save Config* action. Clicking the *Save* icon on the Orchestration Console toolbar does not save the change.

Action	Description
Apply Config	Updates the VM transient configuration. The VM must be running.
Create Template (VM only)	Creates a new template object based on the VM. Other versions of the VM can be cloned from this template. This menu item is replaced by the *Clone* menu item when you right-click a template VM.
Clone (VM template only)	Creates a standalone VM instance from the template but does not power on that instance. This action is unlike the *Provision* action from a template, which performs a *Clone* action, powers on the instance, then destroys the VM instance when it is powered off.
Delete/Destroy Resource	Removes a VM from the *Resources* list in the Orchestration Console. If you want to delete the VM from the host machine, select the *Destroy VM Instance* option. **NOTE:** If you previously added a vDisk to the VM, it is not deleted with the VM unless the vDisk is marked moveable.
Move Disk Images	A “move” is the relocation of VM disk images between two storage devices when the VM is in a not running state (including VMs that are suspended with a checkpoint file). This function does not require shared storage; the move is between separate repositories. Select the storage location from the drop-down menu. You can also move a VM from one VM host machine to another. This is a “cold” migration. VMware Server VMs must be migrated in this manner. If you want to move a VM of considerable size, appropriately increase the timeout fact value in the VM policy. The default value is 2400. For more information on editing the policy, see “Orchestration Provisioning Adapters” in the *NetIQ Cloud Manager 2.3 Component Reference*. If a VM has snapshots, you cannot move the VM but you can register it to a different host if the VM and the host are connected to a shared repository.
Checkpoint	Creates a named snapshot of a VM image. This image is stored on the disk of the repository machine. Xen VMs cannot have a checkpoint applied to them. All the snapshots of a VM are chronologically listed in the `resoruce.vm.snapshots` fact, and the latest snapshot is listed in the `resource.vm.current_snapshot` fact. If the vSphere VM already has snapshots taken through other management consoles, the snapshots are synchronized with the latest snapshot taken through the Orchestration Console, and are listed in the `resource.vm.snapshots` fact.
Restore	Starts a Checkpoint VM (that is, resumes the operations of a VM made into a stored checkpoint from the moment of storage). If the vSphere VM already has snapshots taken through other management consoles, the snapshots are synchronized with the latest snapshot taken through the Orchestration Console, and are listed in the `resource.vm.snapshots` fact.
Remove Template Dependency	Changes a cloned instance of a VM into a VM instance.
Install Agent	Not supported on VMs managed by vSphere.

Action	Description
Personalize	Allows you to customize the VM. This includes changing elements like the DNS server. The changes are made to a VM that is shut down. This action relates to all guest OS-level changes, not hardware changes. With vSphere, guest customization occurs the next time the VM is powered on. The vSphere administrator must have properly configured the vCenter server to handle sysprep and the VMware tools must be installed to the VM. **IMPORTANT:** If you stop or cancel a running Personalize job, the VM is locked and you cannot provision the VM. The VM is automatically released after a period of time.
Shutdown Agent	Shuts down the Orchestration Agent and makes the VM unavailable as a resource.
Create Virtual NIC	Manually creates a vNIC on the VM to configure its network interface configuration by way of a vBridge connection. For more information, see "Creating or Deleting a vNIC in the Orchestration Console" in the *NetIQ Cloud Manager 2.3 Component Reference*. **NOTE:** You can manually delete the new vNIC after it is created by right-clicking the vNIC object in the Explorer tree, then selecting *Delete.*
Create Virtual Disk	Manually creates a vDisk image of a specified size (measured in Mb) to associate to the VM. For more information, see "Creating or Deleting a vDisk in the Orchestration Console" in the *NetIQ Cloud Manager 2.3 Component Reference*. **NOTE:** You can manually delete the new vDisk after it is created by right-clicking the vDisk object in the Explorer tree, then selecting *Delete.*
Cancel Action	Stops an action that has been requested.
Check Host Assignment	Opens a window so you can compare the VM hosts capable of hosting the VM.
Launch Remote Desktop	Launches a VNC terminal in which you can view and control the VM. Specify the credentials configured for the Web service in the appropriate VM policy. Some provisioning adapters require additional setup before you can launch a remote desktop. For more information about this setup for the *vsphere* provisioning adapter, see "Setting Up Orchestration VNC for a VM Managed by vSphere" in the *NetIQ Cloud Manager 2.3 Installation Guide*.

Authorization Constraint Messages Are Not Readily Visible in the Orchestration Console

If you unsuccessfully attempt to provision a VM whose Host/Repository selection has been designated as Automatic, it is possible that a policy with an authorization constraint has been associated to that VM. In this scenario, no message explaining the restriction is displayed.

To confirm that the provisioning has an authorization constraint:

1. In the Explorer tree of the Orchestration Console, select the VM that you want to provision.
2. In the Orchestration Console toolbar, select *Provisioner* to open the provisioning monitor view for that VM.
3. Select the *Show Log* tab to open the provisioning log.

   Scan the log to find errors indicating that the VM could not be provisioned because of authorization constraints in its policy.

7.3.2 Releasing a Virtual Machine from Usage

When the demand and load on your data center decreases, the Orchestration Server analyzes the remaining resources and releases the most appropriate resource. If a VM meets the requirements of the remaining job demands better than a physical machine, the physical machine is released before the VM is released. This dynamic analysis allows you to make sure that the needs of your data center are met.

7.3.3 Managing Virtual Machine Templates

A VM template is a special kind of VM that is not deployed separately.

Review the following tasks to manage VM templates:

- “Making a Virtual Machine Instance into a Template” on page 61
- “Changing a Virtual Machine Template Clone to an Instance” on page 62

Making a Virtual Machine Instance into a Template

1. In the Orchestration Console, right-click the VM.
2. Select *Create Template*.
3. Name the template.
4. Specify a repository.
5. Specify a visible VM host.
6. Select a recommended host for the VMs to be launched on, if any are present.
7. Click *OK*.

When the instance of the template VM is provisioned, it appears as a sub-branch of the template’s location in the resources tree.

This clone functions as an instance of a VM and runs as though it were its own version with its own MAC address and other unique identifiers. The UUID is the only new information that is automatically generated for the clone. All the rest of the new information comes from autoprep, including the MAC

address if an asterisk (*) is placed in the *Mac Address* field in the *Autoprep Network Adapter* section of the *Info/Groups* tab for the template (the default is a blank field, meaning no MAC address is created), and if the *Use Autoprep* check box is enabled in the Create VM from *Template* dialog box.

Changing a Virtual Machine Template Clone to an Instance

1. If you decide to keep a clone VM, go to the Cloud Manager Orchestration Console, right-click it, and select *Remove Template Dependency*.

   The Remove Template Dependency dialog box is displayed.

2. Click *OK*.

NOTE: This procedure works only when a VM instance has been cloned from a template. It does not work when a VM instance is provisioned from a template.

7.3.4 Managed Virtual Machine Actions

You can perform many actions on the VM through the Orchestration Console, from Cloud Manager itself, or you can write jobs to have actions performed on the VMs in your data center. The following table lists the managed VM actions that you can perform or use in a written job.

Table 7-3 *Managed VM Actions*

Action	Description
Provision	Starts a VM. This action clones and start a cloned VM template.
Clone	Creates a new, unique instance of a VM template.
Cold Migrate	Moves the storage location of the configuration and first disk files to another physical storage host. This might allow the VM to start faster.
Shutdown	Stops an active VM instance (including a started template VM). **NOTE:** If you manually shut down a running VM (either by logging in to the VM and shutting it down manually or by using any hypervisor-specific tool to shut down the VM), that VM is still displayed in the Orchestration Console as running. If you then choose the *Shutdown* action rather than the *Resync State* action for that VM, the Orchestration Server job log shows that the shutdown job fails. This condition always exists unless a Resync operation occurs on the VM or when an Orchestration Agent on the VM relays its current state to the Orchestration Server.
Delete/Destroy	Removes a VM from the *Resources* list in the Orchestration Console. If you want to delete the VM from the host machine, select the *Destroy VM Instance* option.
Suspend	Takes a snapshot of an active VM and pauses it in order to move it to another VM host.
Pause	Prevents the VM from obtaining CPU cycles, but it stays resident.
Resume	Allows a paused VM to access the CPU again.
Create Template	Creates a VM template from a VM instance.
Hot Migrate	Changes the association of the VM, which is residing in a shared storage location, from one host machine to another.

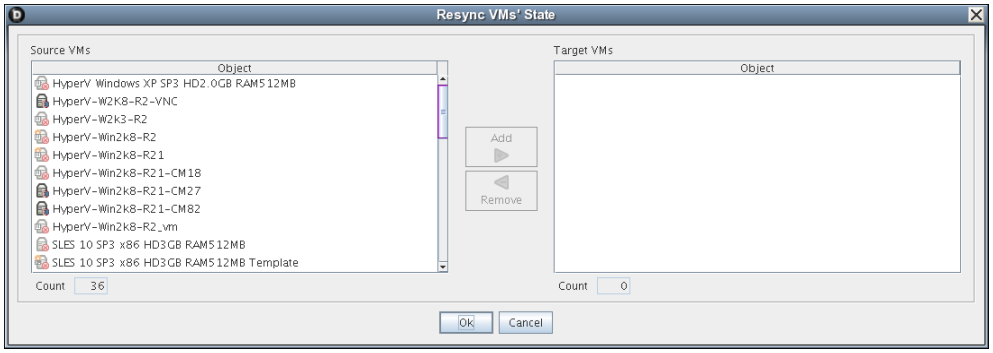
Action	Description
Checkpoint	Create a named snapshot of a moment that can later be accessed to restart from the same point
Restore	Resumes a VM at a previously stored checkpoint.
Install Orchestrator Agent	Opens a VM image and installs the Orchestration Agent.
Make Standalone	Removes the association of a template and makes the active VM into its own instance.
Check Status	Checks the current state of the VM to verify if the VM is provisioned or shut down.
Personalize	Modifies the Orchestration Agent properties and disk image that are currently part of a clone.
Save Config	Transfers changes made to a VM to its permanent image storage.

7.4 Resynchronizing the State of All VMs

To manually verify and ensure that the state of the VMs of all VM hosts displayed in the Orchestration Console is accurate:

1 In the Orchestration Console, click *Provision* > *Resync VM's State*.

The Resync VMs' State dialog box is displayed.



2 In the *Source VMs* pane, select the VMs to be resynchronized, then click *Add*.

The selected VMs are added to the *Target VMs* pane.

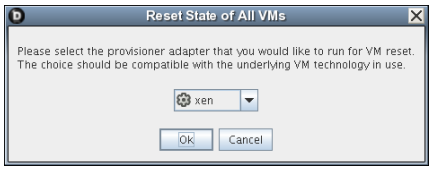
3 Click *OK*.

7.5 Resynchronizing the State of All VMs of a Specific VM Host

To manually verify and ensure that the state of the VMs of a specific VM host displayed in the Orchestration Console is accurate:

1 In the Orchestration Console, click *Provision* > *Reset State of All VMs*.

The Reset State of All VMs dialog box is displayed.



2 Select the VM host whose VMs you want to resynchronize.

3 Click *OK*.

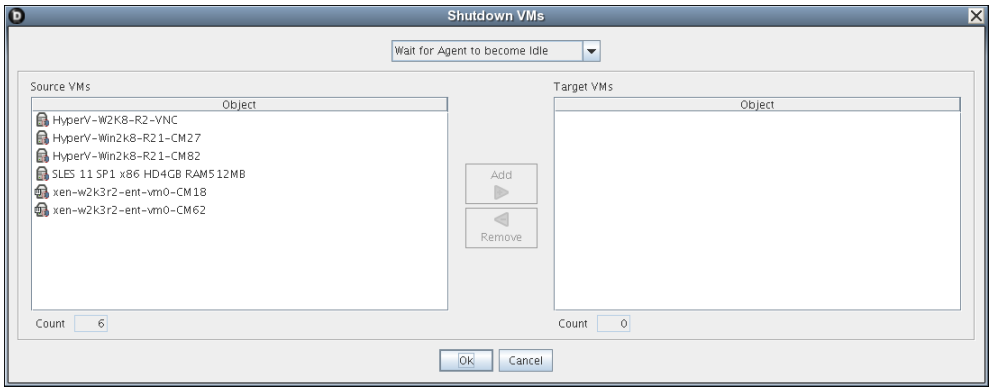
7.6 Shutting Down Multiple VMs

1 In the Orchestration Console, click *Provision* > *Shutdown VMs*.

The Shutdown VMs dialog box is displayed.

2 You can choose to shut down the VMs after the Orchestration Agent becomes idle or to immediately shut down the VMs. By default, the *Wait for Agent to become Idle* option is selected.

3 In the *Source VMs* pane, select the VMs you want to shut down, then click *Add*.



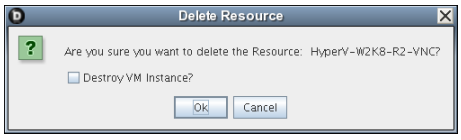
The selected VMs are added to the *Target VMs* pane.

4 Click *OK*.

7.7 Destroying and Deleting a Virtual Machine

1 In the Orchestration Console, right-click the VM in the tree and select *Delete/Destroy Resource*.

The Delete Resource dialog box is displayed.



2 (Optional) To delete a VM from the VM host, select the *Destroy VM Instance* option.

This completely deletes the VM and all its versions from your data center. You cannot restore any version of the VM after you delete it.

If you do not choose this option, the VM is removed from the resource list. However, the actual image of the VM is still stored in its directory.

3 Click *OK*.

If you choose only to delete a VM from your resource tree, you can rediscover the VM by running a discovery job (click *Provision* > *Discover VM Images*).

8 Managing VM Repositories

The Cloud Manager Orchestration Server uses the Repository object to represent where VMs are stored. VMs can be stored on local disks, the datagrid, a network attached storage (NAS), or a storage area network (SAN, Fibre Channel, or iSCSI).

Before VMs can be used by Cloud Manager Orchestration, you must create or discover Repository objects and then discover the VM images within the Repository:

- Section 8.1, “Provisioning a VM from a Local Repository,” on page 67
- Section 8.2, “The Datagrid Repository,” on page 68
- Section 8.3, “Provisioning a VM from a NAS Repository,” on page 68
- Section 8.4, “The SAN Repository,” on page 68

8.1 Provisioning a VM from a Local Repository

By default, the vSphere provisioning adapter creates a local Repository object for local VM images (or datastores) on a VM host when the Orchestration Server performs the Discover VM Hosts action.

NOTE: For vSphere, the Orchestration Server models all repositories the Virtual Center server is aware of. There is not necessarily a local repository for each ESX host, although that is the default.

A local repository represents VMs residing in a VM Host’s local storage where the VMs are only visible to the VM Host. VMs are searched for in the default paths for each adapter.

IMPORTANT: Do not use local repositories for shared directories visible to more than one VM Host. Instead, create a new NAS or SAN repository.

For information on SAN storage, see “The SAN Repository” on page 68.

When discovering VM Images, the adapters use the `location`, `searchpath` and `preferredpath` facts for searching. The `repository.location` is usually the root path, such as `/`.

When the Discover VM Images action is run, the provisioning adapter follows these steps:

- Concatenates the `repository.location` and every element of `repository.searchpath` and searches for VMs in those directories.
- Concatenates the `repository.location` and `repository.preferredpath` and searches for VMs in that directory.

These steps are also followed when searching in NAS and SAN repositories when representing auto-mounted file systems, and when the location, search path, and preferred path are set.

8.2 The Datagrid Repository

By default, a datagrid repository named `zos` is automatically created when the Orchestration Server is initialized. The datagrid repository represents VMs residing in the Orchestration datagrid, which is a storage area on the Orchestration Server.

The `zos` datagrid repository has a location of `grid:///vms`, which points to an area in the datagrid reserved for VM archival storage.

The datagrid repository storage is archival because VMs cannot be run from the datagrid repository. You must move VMs out of the datagrid to a VM Host in order to run them.

8.3 Provisioning a VM from a NAS Repository

The Network Attached Storage (NAS) repository represents VMs stored on a NAS. This is a storage area where VMs are visible to multiple VM hosts, so they can be run by any one of the connected hosts.

For vSphere, you attach the NAS datastore to the host using the vCenter Client, then discover hosts and repositories to bring it into the Cloud Manager Orchestrate Server.

8.4 The SAN Repository

The Storage Area Network (SAN) repository is a single storage server that can be accessed by multiple machines. The Cloud Manager Orchestration Server does not support booting a VM from a SAN repository. SAN repositories can only be used as data disks for VMs.

II Administering the Cloud Manager Application Server and Console

This section includes information that can help you as the Cloud Administrator to administer and set up the Cloud Manager Application Server through its console. It also includes advanced information that you might want to use in special situations.

- Chapter 9, “Setting Up the Cloud Environment,” on page 71
- Chapter 10, “Configuring the Cloud Manager System,” on page 87
- Chapter 11, “Setting Up and Managing Users,” on page 101
- Chapter 12, “Managing the Workload Template Catalog,” on page 113
- Chapter 13, “Managing Organizations,” on page 131
- Chapter 14, “Managing Resources,” on page 141
- Chapter 15, “Managing Business Services,” on page 161
- Chapter 16, “Managing Tasks in the Business Service Workflow,” on page 187
- Chapter 17, “Generating Cloud Manager Reports,” on page 193
- Chapter 18, “Changing Application Server Default Parameters and Values,” on page 195
- Appendix A, “Setting Up Cloud Manager to Log to a Sentinel Collector,” on page 197
- Appendix B, “Integrating Cloud Manager with the Nixu IP Address Management System,” on page 199
- Appendix C, “Using REST APIs to Customize Cloud Manager Behavior,” on page 205

9 Setting Up the Cloud Environment

The following sections provide information to help you set up your Cloud environment. The sections follow the same process and include the same information as the Getting Started view in the Cloud Manager console:

- Section 9.1, “Creating a Zone,” on page 71
- Section 9.2, “Creating Resource Groups Within the Zone,” on page 72
- Section 9.3, “Creating Service Levels for Resource Groups,” on page 73
- Section 9.4, “Creating an Organization,” on page 75
- Section 9.5, “Creating an Organization’s Business Groups,” on page 77
- Section 9.6, “Creating Users and Groups,” on page 78

9.1 Creating a Zone

A Cloud Manager zone is single Cloud Manager Orchestration Server and its managed resources (VM hosts, storage repositories, networks, and so forth). You create a zone by defining a connection from the Cloud Manager Application Server to the Cloud Manager Orchestration Server. After you create the zone, its resources become part of the Cloud environment that you can use to service your customers.

1. On the main navigation bar, click *Getting Started*, then click *Create Zones* (in the *Set Up Your Cloud Environment* list).

   or

   On the main navigation bar, click *Configuration*, click the *Zones* tab, then click *Create*.

2. Provide the following information:

   Name: Provide a unique name for the zone. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Description: If desired, add more information to further identify the zone. The description is displayed in the Cloud Manager console to all users.

   Enabled: Do not change this setting. The zone should be enabled.

   Server Address: Specify the DNS name or IP address of the Orchestration Server.

   Server Port: Specify the port used by the Orchestration Server Web Service.

   Username: Specify the Administrator user name that enables login to the Orchestration Server.

   Password: Specify and confirm the password for the Administrator username.

   Secure Connection: Select this option if the Cloud Manager Application Server is configured for an SSL connection to the Orchestration Server.

3. Click *OK* to create the zone and add it to the list.

For more information about zones, see Chapter 10.1, “Managing Zones,” on page 87.

9.2 Creating Resource Groups Within the Zone

A resource group defines a set of VM hosts that an organization can use for its business services. In addition to the VM hosts, the resource group includes one or more service levels that define the cost of the host resources (vCPUs, memory, storage, and networks) and the service objectives (availability, support response time, and so forth).

- Section 9.2.1, “Resource Group Requirements,” on page 72
- Section 9.2.2, “Shared and Dedicated Resource Groups,” on page 72
- Section 9.2.3, “Creating a Resource Group,” on page 72

9.2.1 Resource Group Requirements

All VM hosts that you include in a resource group must reside in the same Cloud Manager zone. Additionally, the hosts should be identical in terms of hypervisor technology, operating system version, network configuration, storage repository configuration, and hardware capabilities. This ensures a consistent environment for business services regardless of the host. It also ensures that the resource group’s service levels apply to all hosts.

As an example, you might create a Business Critical resource group that consists of high-performance hosts intended for mission-critical applications and services. You assign the resource group a Platinum service level with costs that reflect the more expensive hardware and service contract. Any business service that is provisioned to the resource group also inherits the resource and service costs.

Or, you might create a Lab resource group that consists of standard-performance hosts intended for software testing. You assign the resource group a Bronze service level with costs that reflect the less expensive hardware and service contract.

9.2.2 Shared and Dedicated Resource Groups

A resource group can be shared among multiple organizations, which means that each organization’s business services utilize the same resources, or a resource group can be assigned to only one organization, in which case only that organization’s business services consume the resources.

9.2.3 Creating a Resource Group

1. On the main navigation bar, click *Getting Started*, then click *Create Resource Groups* (in the *Set Up Your Cloud Environment* list).

   or

   On the main navigation bar, click *Resources*, click the *Resource Groups* tab, then click *Create*.
2. If your Cloud Manager system has multiple zones, the Select Zone dialog box is displayed. Select the zone that contains the resources you are grouping, then click *OK* to display the Create Resource Group dialog box.
3. In the *General* fields, provide the following information for the resource group:

   Name: Specify a unique name for the group. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Zone: Displays the zone whose hosts you can add to the group. You cannot change this setting.

Hypervisor: Select the hypervisor technology for the group's hosts. You can add only those hosts that meet the hypervisor criteria.

Workload Repository: The *Default* setting causes a provisioned workload to be stored in the same repository as the VM template used to create it. If you want workloads provisioned to this resource group to be stored in a different shared repository, you must add hosts to the group (see Step 4), then come back and select the shared repository for the workloads. The Workload Repository list is populated only after you add hosts to the resource group.

Group Type: This applies only if VMware vSphere is the selected hypervisor. Select *Host* if you want the resource group to use hosts and host clusters. Select *Resource Pool* if you want the resource group to use a resource pool.

Resource Pool: If you specified *Resource Pool* as the group type, select the resource pool to include in the group.

Description: Provide any additional information for the resource group.

4 If the group type is *Host*, add hosts to the group:

4a Under *Associations*, click the *Hosts* tab.

4b Click *Add* to display the Add Hosts dialog box.

The list displays all available hosts and host clusters in the zone that meet the selected hypervisor criteria. Hosts that are already assigned to another resource group are not displayed.

4c Select the hosts.

You can Shift-click and Ctrl-click to select multiple hosts.

4d Click *OK* to add the selected hosts to the *Hosts* list.

5 Ignore the *Service Levels* tab.

At this point, there are no service levels to assign to the resource group. The next task is to create service levels (see Chapter 9.3, "Creating Service Levels for Resource Groups," on page 73). You assign service levels to resource groups at that time.

6 Ignore the *Networks* tab.

The *Networks* tab shows the networks associated with the hosts you added to the group. The list is view-only so you can't make any changes. However, the list is not generated until you save the resource group. If you want to see the networks at this time, click *Save*, double-click the resource group to open it again, then click the *Networks* tab.

7 Ignore the *Organizations* tab.

At this point, there are no organizations to assign the resource group to. You create organizations and assign resource groups to them later (see Chapter 9.4, "Creating an Organization," on page 75).

8 Click *Save* to add the resource group to the list.

For more information about resource groups, see.Chapter 14, "Managing Resources," on page 141.

9.3 Creating Service Levels for Resource Groups

A resource group has no costs associated with it until you create a service level and assign it to the resource group. The service level defines the cost of the host resources (vCPUs, memory, storage, and networks) and the cost of the service objectives (availability, support response time, and so forth).

You can use the same service level for multiple resource groups. For example, you might have two identical resource groups that require the same service level.

You can also assign multiple service levels to a single resource group. For example, you might create two service levels with the same resource costs but with different service support levels—the first with 24x7x365 support and the second with 12x5x365 support. The user, when requesting a business service, could select the desired service level.

1 On the main navigation bar, click *Getting Started*, then click *Create Service Levels* (in the *Set Up Your Cloud Environment* list).

or

On the main navigation bar, click *Resources*, click *Service Levels*, then click *Create*.

2 In the *General* section, provide the following details for the service level:

Name: Specify a unique name for the service level. This name is displayed in business service workloads.

You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

Creation Date: Displays the current date.

Description: Provide any additional information for the service level.

3 In the *Monthly Resource Costs* fields, define the cost (per month) to use the host resources:

vCPU: Specify the cost per virtual CPU.

Memory: Specify the cost per megabyte (MB) of memory.

Disk: Specify the cost per gigabyte (GB) of disk space.

Network: Specify the cost per network interface card.

4 Assign the service level to the appropriate resource groups:

4a Under *Associations*, click the *Resource Groups* tab.

4b Click *Add* to display the Add Resource Groups dialog box.

4c Select the groups to add.

You can Shift-click and Ctrl-click to select multiple groups.

4d Click *OK*.

5 If you don't want to add service level objectives, click *Save*, then continue with the next task, "Creating Workload Templates" on page 113.

or

If you want to add service level objectives, you must create them first. Click *Save* to save the service level, then do the following:

5a Click the *Service Level Objectives* link.

5b Click *Create* to display the Create Service Level Objective dialog box.

5c Provide the following information:

Name: Specify a name for the objective. This name is displayed in all business service workloads.

You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

Monthly Cost: Specify the cost associated with the objective. If the objective does not have a cost, leave the field empty.

Description: Provide optional text to further identify the service level objective.

Creation Date: Displays the current date.

Objective Type: If this objective represents workload availability, select *Availability*. Otherwise, select *General*.

Value: If the objective type is *Availability*, specify the target availability as a percentage (for example, 99.9). If the objective type is *General*, specify an appropriate objective value.

5d Click *Save*.

6 Add objectives to the service level:

6a Click the *Service Levels* link, select the service level, then click *Edit*.

6b Under *Associations*, click the *Service Level Objectives* tab.

6c Click *Add* to display the Add Service Level Objectives dialog box.

6d Select the objectives to add.

You can Shift-click and Ctrl-click to select multiple objectives.

6e Click *OK* to add the selected objectives to the *Service Level Objectives* list.

6f Click *Save*.

You can include the same objective in more than one service level.

9.4 Creating an Organization

An organization represents a tenant to which you are offering Cloud services. Through the organization, you make resource group assignments that dictate the hosts, service levels, repositories, and networks available to the organization, and make workload template assignments that determine the types of business service workloads available to the organization.

After you create an organization, you can define the organization's membership and assign roles such as Business Group Viewer, Business Service Owner, Business Group Sponsor, Sales Manager, and Organization Manager to those members. Membership and role assignments are covered in the next task, "Creating Users and Groups" on page 78.

1 On the main navigation bar, click *Getting Started*, then click *Create Organizations* (in the *Set Up Your Cloud Environment* list).

or

On the main navigation bar, click *Organizations,* click the *Organizations* tab, then click *Create*.

2 Provide the following details to define the organization:

Name: Specify a unique name for the organization. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

Description: Provide any additional information to identify the organization.

Domains: If you want to automatically create users in this organization, specify the e-mail domains associated with the organization.

Automatic user creation occurs when a valid LDAP user who does not have a Cloud Manager account first logs in. The user's e-mail domain is compared to the e-mail domains defined for the organization. If it matches one of the e-mail domains, the user is added to organization's Members list.

You can associate one or more e-mail domains with the organization. To specify multiple e-mail domains, separate the names with commas (for example, `netiq.com,novell.com,attachmate.com`).

Discount %: If you want a discount applied to all business services created by members of this organization, specify the discount percentage.

Auto Approval: When a user creates a business service request, the request goes through an approval workflow that includes both a Sponsor and an Administrator. The Sponsor is a member of the organization who provides the financial approval for the business service. The Administrator is a System user (such as yourself, another Cloud Administrator, or a Zone Administrator) who provides the resource capacity approval for the business service. You can use Auto Approval to bypass one or both of the approvals.

The organization inherits the Auto Approval settings from the Cloud Manager system settings (accessed through *Configuration* on the main navigation bar). To change the settings for the organization, click *Override*, then configure the settings as desired.

Logo: You can upload a logo file for the organization. Three formats are supported: PNG, JPG, and GIF. Any size is acceptable. Cloud Manager resizes the logo to a maximum of 216x216 pixels, maintaining the width-to-height proportions. For example, a 432x200 image would be resized to 216x100. The logo file is stored on the Cloud Manager Application Server.

To upload a file, mouse over *No Image*, then click *Upload New Image*. Browse for and select the image, then click *OK* to upload it to the Cloud Manager Application Server.

3 Add the resource groups that you want the organization to have access to:

3a Under *Membership and Access*, click the *Resource Groups* tab.

3b Click *Add* to display the Add Resource Groups dialog box.

3c Select the resource groups you want to add.

You can Shift-click and Ctrl-click to select multiple groups.

3d Click *OK* to add the selected resource groups to the *Resource Groups* list.

4 Add the workload templates that you want the organization to have access to:

4a Under *Membership and Access*, click the *Workload Templates* tab.

4b Click *Add* to display the Add Workload Templates dialog box.

4c Select the workload templates.

You can Shift-click and Ctrl-click to select multiple workload templates.

4d Click *OK* to add the selected workload templates to the *Workload Templates* list.

5 Add the networks that you want the organization to have access to.

The available networks are determined by the VM hosts included in the resources group. However, to enable you to provide isolated networks for organizations that share the same resource group, the networks from a resource group are not automatically assigned to an organization when you add the resource group. Instead, you must separately add the networks you want assigned to the organization.

5a Under *Membership and Access*, click the *Networks* tab.

5b Click *Add* to display the Add Networks dialog box.

5c Select the networks.

You can Shift-click and Ctrl-click to select multiple workload templates.

5d Click *OK* to add the selected networks to the *Networks* list.

6 Ignore the *Users* tab and the *Business Groups* tab at this time.

The *Users* tab lets you add members to the organization and assign them roles within the organization. The *Business Groups* tab lets you view the sub-units that have been created for the organization. Creating business groups is covered in the next task, “Creating an Organization’s Business Groups” on page 77. Creating users is covered in Chapter 9.6, “Creating Users and Groups,” on page 78.

7 Click *Save* to create the organization and add it to the list.

For more information about organizations, see Chapter 13, “Managing Organizations,” on page 131.

9.5 Creating an Organization’s Business Groups

An organization includes one or more business groups. A business group represents a unit within the organization, such as a department or cost center, for which business services can be deployed.

A business group can be assigned access to all of an organization’s resources or only some of the resources. When a business service is created for a business group, it uses only the assigned resources. Multiple business groups can be assigned the same resources, which means that the resources become shared resources.

1 On the main navigation bar, click *Getting Started*, then click *Create Business Groups* (in the *Set Up Your Cloud Environment* list).

or

On the main navigation bar, click *Organizations,* click the *Business Groups* tab, then click *Create*.

2 Provide the following details to define the business group:

Name: Specify a unique name for the group. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

Organization: Select the organization where you want to add the business group.

Description: Provide any additional information to identify the business group.

Auto Approval: Business service requests require approval from both the group’s Sponsor and an administrator. The *Auto Approval* settings let you bypass one or both of the approvers.

The business group inherits the *Auto Approval* settings from its organization. To change the settings for the business group, click *Override*, then configure the settings as desired.

Costs: The business group inherits the *Costs* setting from its organization. To change the setting for the business group, click *Override*, then configure the setting as desired. The *Show* setting allows group members to see cost information for workloads. The *Hide* setting prevents group members from seeing cost information.

3 Add the resource groups that you want the business group to have access to:

3a Under *Membership and Access*, click the *Resource Groups* tab.

3b Click *Add* to display the Add Resource Groups dialog box.

The list displays the organization’s resource groups. A business group is limited to the resource groups assigned to its organization.

3c Select the resource groups you want to add.

You can Shift-click and Ctrl-click to select multiple groups.

3d Click *OK* to add the selected resource groups to the *Resource Groups* list.

4 Add the workload templates that you want the business group to have access to:

4a Under *Membership and Access*, click the *Workload Templates* tab.

4b Click *Add* to display the Add Workload Templates dialog box.

The list displays the organization's workload templates. A business group is limited to the workload templates assigned to its organization.

4c Select the workload templates.

You can Shift-click and Ctrl-click to select multiple workload templates.

4d Click *OK* to add the selected workload templates to the *Workload Templates* list.

5 Add the networks that you want the business group to have access to.

The available networks are determined by the VM hosts included in the resources groups you added in Step 3. However, to enable you to provide isolated networks for business groups that share the same resource group, the networks from a resource group are not automatically assigned to a business group when you add the resource group. Instead, you must separately add the networks you want assigned to the business group.

5a Under *Membership and Access*, click the *Networks* tab.

5b Click *Add* to display the Add Networks dialog box.

The list displays the organization's networks. A business group is limited to the networks assigned to its organization.

5c Select the networks.

You can Shift-click and Ctrl-click to select multiple networks.

5d Click *OK* to add the selected networks to the *Networks* list.

6 Ignore the *Users* tab and the *Business Services* tab at this time.

The *Users* tab lets you add members to the business group and assign them roles within the business group. The *Business Services* tab lets you view the business services that are currently deployed for the business unit. Creating users for the business group is covered in the next task, "Creating Users and Groups" on page 78.

7 Click *Save* to create the business group and add it to the list.

For more information about business groups, see Chapter 13, "Managing Organizations," on page 131.

9.6 Creating Users and Groups

Access to Cloud Manager requires a Cloud Manager user account. Through the account, a user receives rights to perform various roles in the Cloud Manager system, in an organization, or in both. Rights can also be assigned to user groups to enable all members of the group to perform specific roles.

You can create users and groups by manually entering information or by importing information from your LDAP authentication source.

- Section 9.6.1, "Manually Creating Users," on page 79
- Section 9.6.2, "Manually Creating User Groups," on page 80
- Section 9.6.3, "Importing Users from LDAP," on page 82
- Section 9.6.4, "Importing User Groups from LDAP," on page 84

9.6.1 Manually Creating Users

1 On the main navigation bar, click *Getting Started*, then click *Create Users and Groups* (in the *Set Up Your Cloud Environment* list).

or

On the main navigation bar, click *Users*, then click the *Users* tab.

2 On the *Users* tab, click *Create* to display the Create User dialog box.

3 Provide the following details to define the user:

Full Name: Specify the user's full name as you want it to appear in NetIQ Cloud Manager.

E-Mail Address: Specify the user's e-mail address as defined in their LDAP authentication account. If necessary, you can specify more than one address; use commas to separate addresses.

The e-mail address enables the Cloud Manager system to send messages (tasks, notifications, and so forth) to the user as needed.

Phone Number: This field is optional. Specify a contact number if desired.

4 Select the user's scope:

Organization: An organization scope enables the user to perform roles within a specific organization. The roles are Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor.

To give the user an organization scope, select *Organization*, then select the organization in which to place the user.

System: A system scope enables the user to administer the Cloud Manager system. The roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. In addition, a System user can be given any of the organization roles.

NOTE: System scope also implies the Zone Administrator role, though it is not explicitly listed. Instead, specific zones are associated to these users, making the Zone Administrator role implicit.

---

5 (Organization user only) If you want the user to always be able to view business service costs regardless of the *Costs* setting for a business group, select *Always show costs*.

An organization's or business group's *Costs* setting can be set to *Show* or *Hide*. The purpose of the *Always show costs* setting is to ensure that business service costs are always visible to the user even if the *Costs* setting is set to *Hide*.

For example, you might want to select this option for users who are Sponsors. This ensures that the users can always see costs even if the organization or business group is set to hide costs.

6 (System user only) Assign system-level roles to the user.

The system-level roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. These roles can be assigned only to System users.

NOTE: System scope also implies the Zone Administrator role, though it is not explicitly listed. Instead, specific zones are associated to these users, making the Zone Administrator role implicit.

6a To assign the Approver, Build Administrator, Catalog Manager, or Cloud Administrator role, click the *System* tab, click *Add*, select the desired roles, then click *OK*.

6b To assign the Zone Administrator role, click the *Zone* tab, click *Add*, select the desired zone, then click *OK*.

7 Assign organization-level roles to the user.

The organization-level roles are Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor. The Approver, System View, and Build Administrator roles can be assigned only to System users. The Sales Manager role can be assigned only to Organization users. The other roles can be assigned to both System users and Organization users.

Several of the roles can be assigned at the organization, business group, or business service level. For example, you can make a user a Sponsor for a business group, in which case the user can approve requests for business services from that business group only. Or, you can make the user a Sponsor for the organization, in which case the user can approve requests for all business services in the organization.

7a Click the *Organization* tab to add a role at the organization level, click the *Business Group* tab to add a role at the business group level, or click the *Business Service* tab to add a role at the business service level.

7b Click the role that you want to assign

For example, if you selected the *Business Group* tab and you want to enable the user to create business services for the business group, click *Business Service Owner*.

7c Click *Add*, select the object (organization, business group, or business service) to which you want the role to apply, then click *OK* to add it to the list.

8 Ignore the *Membership* tab at this time.

The *Membership* tab lets you add users to groups. You must create the groups first. This task is discussed in "Manually Creating User Groups" on page 80 and "Importing User Groups from LDAP" on page 84

9 When you have finished assigning roles to the user, click *Save*.

For more information about users and roles, see Chapter 11, "Setting Up and Managing Users," on page 101.

9.6.2 Manually Creating User Groups

Rather than assign roles to individual users, you can create user groups and assign roles to the user groups. Users (and other user groups) that are added to a group inherit the group's roles.

User group roles are cumulative. If you add a user to a group, the user retains its directly assigned roles and also gains the roles inherited from the group.

1 On the main navigation bar, click *Getting Started*, then click *Create Users and Groups* (in the *Set Up Your Cloud Environment* list).

or

On the main navigation bar, click *Users*.

2 Click the *User Groups* tab, then click *Create* to display the Create User Group dialog box.

3 Provide the following details to define the user group:

Full Name: Specify the group's full name as you want it to appear in NetIQ Cloud Manager.

E-Mail Address: This field is optional. If you enter an e-mail address, any messages generated for the group's roles are sent to the e-mail address. If you don't enter an e-mail address, the messages are sent to the group members' addresses.

4 Select the group's scope:

Organization: An organization scope enables the group to be assigned roles within a specific organization. The roles are Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor.

To give the group an organization scope, select *Organization*, then select the organization in which to place the group.

System: A system scope enables the group to be assigned roles for the Cloud Manager system. The roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. In addition, a System group can be given any of the organization roles.

NOTE: System scope also implies the Zone Administrator role, though it is not explicitly listed. Instead, specific zones are associated to these users, making the Zone Administrator role implicit.

---

5 (System user groups only) Assign system-level roles to the group.

The system-level roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. These roles can be assigned only to System user groups.

5a To assign the Approver, Build Administrator, Catalog Manager, or Cloud Administrator role, click the *System* tab, click *Add*, select the desired roles, then click *OK*.

5b To assign the Zone Administrator role, click the *Zone* tab, click *Add*, select the desired zone, then click *OK*.

6 Assign organization-level roles to the group.

The organization-level roles are Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor. The Approver and Build Administrator roles can be assigned only to System user groups. The other roles can be assigned to both System and Organization user groups.

Several of the roles can be assigned at the organization, business group, or business service level. For example, you can make a user group a Sponsor for a business group, in which case the group members can approve requests for business services from that business group only. Or, you can make the user group a Sponsor for the organization, in which case the group members can approve requests for all business services in the organization.

6a Click the *Organization* tab to add a role at the organization level, click the *Business Group* tab to add a role at the business group level, or click the *Business Service* tab to add a role at the business service level.

6b Click the role that you want to assign.

For example, if you selected the *Business Group* tab and you want to enable the user group to create business services for the business group, click *Business Service Owner*.

6c Click *Add*, select the object (organization, business group, or business service) to which you want the role to apply, then click *OK* to add it to the list.

7 Add members to the group:

7a Click the *Membership* tab.

7b Click *Members*, then click *Add* to display the Add Members dialog box.

7c Select the users and user groups you want to add to the group.

You can Shift-click and Ctrl-click to select multiple users and groups.

7d Click *OK* to add the users and user groups to the Members list.

8 When you have finished assigning roles and adding members, click *Save*.

For more information about user groups and roles, see Chapter 11, “Setting Up and Managing Users,” on page 101.

9.6.3 Importing Users from LDAP

You can create users by importing information from your LDAP authentication source. You can import users as System or Organization users. After you import a user, you can assign roles to the user.

1. On the main navigation bar, click *Getting Started*, then click *Create Users and Groups* (in the *Set Up Your Cloud Environment* list).

   or

   On the main navigation bar, click *Organizations*.

2. If you want to import Organization users, click the *Organizations* tab, select the target organization for the import, click *Edit* to display the Edit Organization dialog box, then click *Import* (located above the *Members* list on the *Users* tab).

   or

   If you want to import System users, click *Configuration* (on the main navigation bar) to display the System Configuration dialog box, click *System Users*, click the *Members* tab, then click *Import*.

3. Authenticate to the LDAP directory:

   3a Click the *LDAP* tab.

   3b In the *LDAP Location* section, fill in the following fields:

   Host: Specify the FQDN (fully qualified domain name) or IP address of the host machine running the LDAP server. For example, `ldap.mycompany.com` or `123.45.67.8`.

   Port: Specify the TCP port (on the host machine) where the LDAP server is listening for LDAP connections. The standard port for non-SSL connections is 389. The standard port for SSL connections is 636.

   Use SSL: If the Cloud Manager Application Server is configured for an SSL connection to the LDAP server, select this option to enable the secure connection.

   3c In the *Search Bind Account* section, fill in the following fields:

   DN: Specify an account that has search rights to the directory location from which you want to import users. For example, `cn=Administrator,cn=Users,dc=MyCompany,dc=com`

   Password: Specify the password for the account.

   Password Confirm: Confirm the password for the account.

   3d Click *Test Connection*.

   If the connection is successful, the Test Status is displayed as *Passed*. If the connection is not successful, validate the connection information and try again.

4. Import users:

   4a Click the *Import* tab.

   4b Click *Add*.

   When you click *Add*, an new import entry is added to the list. You use the fields below the list to define the entry.

   4c In the DN field, use standard LDAP notation (`ou=provo,dc=netiq,dc=com`) to specify the distinguished name for the target container or object, then click *Validate*.

   If you specify a container, all users located within the container are imported. If you only want to import one user, specify the DN of the user object.

   4d If you specified a container for import, select *Users*.

4e If you specified a container for import, select *Scan Tree* if you want to import users located in its subcontainers.

4f Click *Import*.

The imported users are added to the *Members* list. Users are identified by the icon.

5 When you have finished importing users, click *OK* or *Save* to close the dialog box.

6 Assign roles to the users:

6a On the main navigation bar, click *Users*.

6b Click the *Users* tab, select the user to whom you want to assign roles, then click *Edit*.

6c (System user only) Assign system-level roles.

The system-level roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. These roles can be assigned only to System users.

NOTE: System scope also implies the Zone Administrator role, though it is not explicitly listed. Instead, specific zones are associated to these users, making the Zone Administrator role implicit.

6c1 To assign the Approver, Build Administrator, Catalog Manager, Cloud Administrator, or System View role, click the *System* tab, click *Add*, select the desired roles, then click *OK*.

6c2 To assign the Zone Administrator role, click the *Zone* tab, click *Add*, select the desired zone, then click *OK*.

6d Assign organization-level roles.

The organization-level roles are Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor. The Approver and Build Administrator roles can be assigned only to System users. The Sales Manager role can be assigned only to Organization users. The other roles can be assigned to both System users and Organization users.

Several of the roles can be assigned at the organization, business group, or business service level. For example, you can make a user a Sponsor for a business group, in which case the user can approve requests for business services from that business group only. Or, you can make the user a Sponsor for the organization, in which case the user can approve requests for all business services in the organization.

6d1 Click the *Organization* tab to add a role at the organization level, click the *Business Group* tab to add a role at the business group level, or click the *Business Service* tab to add a role at the business service level.

6d2 Click the role that you want to assign

For example, if you selected the *Business Group* tab and you want to enable the user to create business services for the business group, click *Business Service Owner*.

6d3 Click *Add*, select the object (organization, business group, or business service) to which you want the role to apply, then click *OK* to add it to the list.

6e When you have finished assigning roles to the user, click *Save*.

For more information about users and roles, see Chapter 11, “Setting Up and Managing Users,” on page 101.

9.6.4 Importing User Groups from LDAP

You can create user groups by importing them from your LDAP authentication source. After you import a group, you can assign roles to the group.

An imported user group's membership is maintained in the LDAP authentication source. Any users who are members of the user group in the LDAP source receive the roles that are assigned to the user group in Cloud Manager.

An imported user group's members are not imported and do not display in the group's *Members* list. In addition, you cannot manually add users or user groups to an imported group.

1. On the main navigation bar, click *Getting Started*, then click *Create Users and Groups* (in the *Set Up Your Cloud Environment* list).

   or

   On the main navigation bar, click *Organizations*.

2. If you want to import Organization user groups, click the *Organizations* tab, select the target organization for the import, click *Edit* to display the Edit Organization dialog box, then click *Import* (located above the *Members* list on the *Users* tab).

   or

   If you want to import System user groups, click *Configuration* (on the main navigation bar) to display the System Configuration dialog box, click *System Users*, click the *Members* tab, then click *Import*.

3. Authenticate to the LDAP directory:

   3a Click the *LDAP* tab.

   3b In the *LDAP Location* section, fill in the following fields:

   Host: Specify the FQDN (fully qualified domain name) or IP address of the host machine running the LDAP server. For example, `ldap.mycompany.com` or `123.45.67.8`.

   Port: Specify the TCP port (on the host machine) where the LDAP server is listening for LDAP connections. The standard port for non-SSL connections is 389. The standard port for SSL connections is 636.

   Use SSL: If the Cloud Manager Application Server is configured for an SSL connection to the LDAP server, select this option to enable the secure connection.

   3c In the *Search Bind Account* section, fill in the following fields:

   User DN: Specify an account that has read rights to the directory location from which you want to import users. For example, `cn=Administrator,cn=Users,dc=MyCompany,dc=com`

   Password: Specify the password for the account.

   Password Confirm: Confirm the password for the account.

   3d Click *Test Connection*.

   If the connection is successful, the Test Status is displayed as *Passed*. If the connection is not successful, validate the connection information try again.

4. Import user groups:

   4a Click the *Import* tab.

   4b Click *Add*.

   When you click *Add*, an new import entry is added to the list. You use the fields below the list to define the entry.

4c In the DN field, use standard LDAP notation (`ou=provo,dc=netiq,dc=com`) to specify the distinguished name for the target container or object, then click *Validate*.

If you specify a container, all user groups located within the container are imported. If you only want to import one user group, specify the DN of the user group object.

4d If you specified a container for import, select *Groups*.

4e If you specified a container for import, select *Scan Tree* if you want to import user groups located in its subcontainers.

4f Click *Import*.

The imported user groups are added to the *Members* list. User groups are identified by the  icon.

5 When you have finished importing user groups, click *OK* or *Save* to close the dialog box.

6 Assign roles to the groups:

6a On the main navigation bar, click  *Users*.

6b Click the *User Groups* tab, select the user group to which you want to assign roles, then click *Edit*.

6c (System user groups only) Assign system-level roles.

The system-level roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. These roles can be assigned only to System user groups.

NOTE: System scope also implies the Zone Administrator role, though it is not explicitly listed. Instead, specific zones are associated to these user groups, making the Zone Administrator role implicit.

6c1 To assign the Approver, Build Administrator, Catalog Manager, System View, or Cloud Administrator role, click the *System* tab, click *Add*, select the desired roles, then click *OK*.

6c2 To assign the Zone Administrator role, click the *Zone* tab, click *Add*, select the desired zone, then click *OK*.

6d Assign organization-level roles.

The organization-level roles are Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor. The Sales Manager can be assigned only to Organization user groups. The Approver and Build Administrator roles can be assigned only to System user groups. The other roles can be assigned to both System and Organization user groups.

Several of the roles can be assigned at the organization, business group, or business service level. For example, you can make a user group a Sponsor for a business group, in which case the group members can approve requests for business services from that business group only. Or, you can make the user group a Sponsor for the organization, in which case the group members can approve requests for all business services in the organization.

6d1 Click the *Organization* tab to add a role at the organization level, click the *Business Group* tab to add a role at the business group level, or click the *Business Service* tab to add a role at the business service level.

6d2 Click the role that you want to assign

For example, if you selected the *Business Group* tab and you want to enable the user group to create business services for the business group, click *Business Service Owner*.

6d3 Click *Add*, select the object (organization, business group, or business service) to which you want the role to apply, then click *OK* to add it to the list.

6e When you have finished assigning roles to the user group, click *Save*.

For more information about user groups and roles, see Chapter 11, “Setting Up and Managing Users,” on page 101.

10 Configuring the Cloud Manager System

The following sections provide information to help configure and manage your Cloud Manager system:

- Section 10.1, “Managing Zones,” on page 87
- Section 10.2, “Customizing the Capacity Thresholds and Data Refresh Interval,” on page 89
- Section 10.3, “Managing System Users, User Groups, and Roles,” on page 90
- Section 10.4, “Configuring Email Notifications,” on page 90
- Section 10.5, “Configuring Remote Console Access to Workloads,” on page 91
- Section 10.6, “Configuring Auto Approval for Business Service Requests,” on page 94
- Section 10.7, “Customizing the Cost Update Schedule,” on page 94
- Section 10.8, “Customizing the Cloud Manager Console Interface,” on page 95
- Section 10.9, “Setting Up the Web Application Launcher (Launch Pad),” on page 96

10.1 Managing Zones

A Cloud Manager zone is single Cloud Manager Orchestration Server and its managed resources (VM hosts, storage repositories, networks, and so forth). The following sections provide instructions to help you manage your zones:

- Section 10.1.1, “Creating Zones,” on page 87
- Section 10.1.2, “Disabling Zones,” on page 88
- Section 10.1.3, “Enabling Zones,” on page 88
- Section 10.1.4, “Removing Zones,” on page 89

10.1.1 Creating Zones

Roles that Can Perform This Task: Cloud Administrator

A Cloud Manager zone is a single Cloud Manager Orchestration Server and its managed resources (VM hosts, storage repositories, networks, and so forth). You create a zone by defining a connection to the Orchestration Server. After you create the zone, its resources become part of the Cloud environment that you can use to service your customers.

1. On the main navigation bar, click *Configuration*, click the *Zones* tab, then click *Create*.
2. Provide the following information:

   Name: Provide a unique name for the zone. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Description: If desired, add more information to further identify the zone. The description is displayed in Cloud Manager to administrators only.

Enabled: Leave this setting selected.

Server Address: Specify the DNS name or IP address of the Orchestration Server.

Server Port: Specify the port used by the Orchestration Server Web Service.

Username: Specify the Administrator user name that enables login to the Orchestration Server.

Password: Specify and confirm the password for the user name you supplied.

Secure Connection: Select this option if the Cloud Manager Application Server is configured for an SSL connection to the Orchestration Server.

3 Click *OK* to create the zone and add it to the list.

4 When you have finished creating zones, click *OK* to close the System Configuration dialog box.

10.1.2 Disabling Zones

Roles that Can Perform This Task: Cloud Administrator

The primary purpose for disabling a zone is to perform maintenance tasks on the Cloud Manager Orchestration Server without receiving error messages and exceptions in the Cloud Manager console because the server is down. While the zone is disabled, the following occurs:

- Any objects (hosts, VM templates, and so forth) provided by the zone's Orchestration Server are no longer displayed in the Cloud Manager console. However, Cloud Manager objects (resource groups, service levels, workload templates, and so forth) are still displayed. A user receives a "Zone Disabled" error when adding, editing, or deleting any Cloud Manager object that includes an Orchestration Server object.
- Any workloads deployed in the zone continue in their current state (unless an issue with the zone causes them to be stopped). In the Cloud Manager console, a workload's state displays as "Unknown" and the owner cannot cycle the workload.

To disable a zone:

1 On the main navigation bar, click *Configuration*, then click the *Zones* tab.

2 Select the zone to disable, then click *Edit*.

3 Deselect the *Enabled* check box.

4 Click *OK* to disable the zone.

10.1.3 Enabling Zones

Roles that Can Perform This Task: Cloud Administrator

To enable a zone:

1 On the main navigation bar, click *Configuration*, then click the *Zones* tab.

2 Select the zone to enable, then click *Edit*.

3 Select the *Enabled* check box.

4 Provide the Administration password for the Cloud Manager Orchestration Server.

5 Click *OK* to enable the zone.

10.1.4 Removing Zones

Roles that Can Perform This Task: Cloud Administrator

You cannot remove a zone that has 1) business service workloads running in the zone, 2) business service requests in process, or 3) resource groups, service levels, and workload templates associated with the zone.

To remove a zone:

1. On the main navigation bar, click *Resources*, then click the *Zones* tab.

   or

   On the main navigation bar, click *Users*, then click either the *Users* or *User Groups* tab.
2. Select the zone to remove, then click *Remove*.

   If the *Remove* action is not available, the zone has workload templates or resource groups associated with it.
3. Click *Yes* to confirm removal of the zone.

10.2 Customizing the Capacity Thresholds and Data Refresh Interval

Roles that Can Perform This Task: Cloud Administrator

The Capacity view provides resource capacity and usage information for the organizations and zones in your system.

The Capacity view has two settings you can customize: *Thresholds* and *Data Refresh Interval*. *Thresholds* lets you set the Warning threshold and Problem threshold for used resource capacity. *Data Refresh Interval* lets you determine how often capacity data is collected.

1. On the main navigation bar, click *Configuration*, then click *Capacity*.
2. Under *Thresholds*, select the Warning and Problem thresholds you want for each resource (Memory, CPU, and Storage).

   The Capacity view displays resource usage as a percentage of the total capacity. Visually, this is displayed as a horizontal bar ranging from 0 percent to 100 percent.

   The Warning and Problem thresholds help you quickly see when resource usage is approaching your resource capacity. Up to the Warning threshold, usage is displayed as a green bar. When the Warning threshold is reached, usage is displayed as a yellow bar. When the Problem threshold is reached, usage is displayed as a red bar.
3. Under *Data Refresh Interval*, specify a frequency for collecting capacity data from your system.
   - **3a** (Optional) Select *No Refresh* if you choose not to periodically refresh the data. This option keeps the original capacity data for display on the Capacity dashboard.
   - **3b** (Optional) Select *Interval* to specify how frequently you want to initiate capacity data collection.

Keep in mind that data collection can be an intensive process. The frequency of data collection and the duration of data collection (which is dependent on the size of your system) could affect your system's performance. You should specify an interval that balances the need for up-to-date information against the potential impact to system performance while the Capacity engine is running.

The system must update the capacity data at least once before you can accurately estimate the interval. You can base your estimate on the data displayed in the Capacity Update dialog box when you select *Update* in the Capacity view.

3c (Optional) Select *Daily Time* to specify the time each day (based on the Cloud Manager Application Server's clock and time zone) when you want Cloud Manager to collect capacity data.

4 Click *OK* to save your change.

10.3 Managing System Users, User Groups, and Roles

Roles that Can Perform This Task: Cloud Administrator

You can add users and user groups to the system and then assign roles to the users and groups so that they can perform specific functions within the system or within organizations.

Users, user groups, and roles are covered in Chapter 11, "Setting Up and Managing Users," on page 101.

10.4 Configuring Email Notifications

Roles that Can Perform This Task: Cloud Administrator

Cloud Manager can send e-mail messages to remind task owners about tasks that need to be completed and to notify Business Service Owners of business services that are about to expire or that have expired. For Cloud Manager to do this, you must provide the connection information for an SMTP server to route the messages. You can also customize the schedule for the message notifications.

1 On the main navigation bar, click *Configuration*, then click *Email Configuration*.

2 Configure the SMTP server connection:

Host: Specify the IP address or DNS name of the host running the SMTP server.

Server Port: Specify the port on which the SMTP server listens for incoming messages.

From Address: Specify the no-reply address you want to use as the sender of the messages.

3 (Conditional) If the SMTP server requires an authentication user name and password, select this option, then specify the user name and password.

4 Click *OK*.

10.5 Configuring Remote Console Access to Workloads

The Cloud Manager console provides remote console access to business service workloads via an embedded Flash VNC application. The application can connect to workloads either directly or through a VNC repeater (proxy).

By default, the Cloud Manager console is configured to use the VNC repeater included with the Cloud Manager Application Server. Alternately, you can set up an external VNC repeater or configure the VNC application to connect directly to workloads. Each solution has advantages and disadvantages

Solution	Advantages	Disadvantages
Built-In Repeater	◆ Minimal setup ◆ Supports NAT and firewalls ◆ If used with NAT or a firewall, use can be limited to Cloud Manager users	◆ VNC traffic flows through Cloud Manager Application Server, increasing workload on a single server
External Repeater	◆ Supports NAT and firewalls ◆ Offloads VNC traffic from the Cloud Manager Application Server ◆ Scalable by clustering the repeater	◆ Increased setup ◆ VNC requests are not authenticated through Cloud Manager
Direct Connection	◆ Most scalable	◆ Each workload must include a VNC server ◆ No support for NAT or firewalls ◆ VNC requests are not authenticated through Cloud Manager

The following sections provide instructions for configuring each of the remote console access solutions:

- Section 10.5.1, “Disabling Remote Console Access,” on page 91
- Section 10.5.2, “Setting Up the Built-In VNC Repeater,” on page 92
- Section 10.5.3, “Setting Up an External VNC Repeater,” on page 92
- Section 10.5.4, “Setting Up Direct Connections,” on page 92
- Section 10.5.5, “Enabling Repeater SSL Encryption,” on page 93
- Section 10.5.6, “Enabling File Uploading,” on page 93

10.5.1 Disabling Remote Console Access

Roles that Can Perform This Task: Cloud Administrator

If you don’t want users to be able to access workloads through the Cloud Manager console, you can disable remote console access. This disables remote console access to all workloads through the Cloud Manager console only. It does not disable VNC on the host or the workload.

1. On the main navigation bar, click *Configuration*, then click *Remote Console*.
2. In the *Connection* field, select *Disable*.
3. Click *OK*.

10.5.2 Setting Up the Built-In VNC Repeater

Roles that Can Perform This Task: Cloud Administrator

To have the Cloud Manager console use the built-in VNC repeater:

1. On the main navigation bar, click *Configuration*, then click *Remote Console*.
2. In the *Connection* field, select *Use built-in VNC repeater*.
3. If the VNC repeater requires a static port for reasons such as firewall support, specify the port in the *Repeater Port* field. Otherwise, leave the field blank so that the VNC repeater dynamically selects an available port when it starts.
4. Click *OK*.

10.5.3 Setting Up an External VNC Repeater

Roles that Can Perform This Task: Cloud Administrator

To have the Cloud Manager console use an external VNC repeater:

1. Install the VNC repeater by using the product’s documentation.
2. Configure the repeater to respond to both Flash policy requests and VNC proxy requests.
3. In the Cloud Manager console, configure the remote console to use the external repeater:
   - **3a** On the main navigation bar, click *Configuration*, then click *Remote Console*.
   - **3b** In the *Connection* field, select *Use external VNC repeater*.
   - **3c** In the *Repeater Address* field, specify the DNS or IP address of the VNC repeater’s server.
   - **3d** In the *Repeater Port* field, specify the port assigned to the repeater.
   - **3e** Click *OK*.

10.5.4 Setting Up Direct Connections

Roles that Can Perform This Task: Cloud Administrator

To have the Cloud Manager console connect directly to workloads:

1. Make sure that each VM host or VM is configured with a VNC Server.

Depending on the hypervisor, the VM host might handle the VNC requests for the VM or the VM might handle the requests. Refer to your hypervisor documentation for information about how your hypervisor handles VNC requests to VMs.

2 Configure the VNC Server to respond to Flash policy requests.

3 In the Cloud Manager console, configure the remote console to use a direct connection:

3a On the main navigation bar, click *Configuration*, then click *Remote Console*.

3b In the *Connection* field, select *Connect directly*.

3c Click *OK*.

10.5.5 Enabling Repeater SSL Encryption

Roles that Can Perform This Task: Cloud Administrator

To enable SSL encryption of VNC traffic between your browser and the VNC repeater (making it difficult for an outside entity to intercept and analyze activity between your browser and the repeater):

1 On the main navigation bar, click *Configuration*, then click *Remote Console*.

2 Select the *Enable Repeater SSL Encryption* check box.

3 In the *Repeater Keystore* field, enter the path to a Java keystore where the SSL key to the VNC Repeater is stored.

By default, this field is populated from the original Cloud Manager SSL configuration (if that option was chosen).

4 In the *Keystore Password* field, enter the password to the Java keystore where the SSL key to the VNC Repeater is stored.

By default, this field is populated from the original Cloud Manager SSL configuration (if that option was chosen).

5 (Conditional) If you select the built-in repeater, specify the path to the repeater keystore and passwords. The first password (required) opens the keystore file. The second password (optional) retrieves the private key within the file. The need for the second password depends on the settings you used when you generated the keystore.

NOTE: The fields on this page validate the keystore and passwords as you make changes: if you enter an incorrect password, the field displays a red asterisk.

6 Click *OK*.

10.5.6 Enabling File Uploading

Roles that Can Perform This Task: Cloud Administrator

Cloud Manager users can upload files to their workloads while attached by using the Remote Console window. Uploaded files appear as a mounted CD-ROM from within the virtual machine. This feature is only available for workloads that have a CD-ROM device defined.

1. On the main navigation bar, click *Configuration*, then click *Remote Console*.
2. Select the *Enable File Uploads* check box.
3. Click *OK*.

10.6 Configuring Auto Approval for Business Service Requests

Roles that Can Perform This Task: Cloud Administrator

When a business service request is submitted, the request goes through an approval workflow that requires both a Sponsor approval and an Administrator approval. The Sponsor approval is intended to be a financial check and the Administrator approval is intended to be a resource capacity check.

You can enable automatic Sponsor approval, Administrator approval, or both for your system. This eliminates the need to assign users as Sponsors and Approvers for organizations or business group. If you don't want the system settings to apply to an organization or a business group, you can override the settings at the organization or business group.

To configure the system Auto Approval settings:

1. On the main navigation bar, click *Configuration*, then click *Tasks*.
2. Select *Sponsor* to enable automatic Sponsor approval.
3. Select *Administrator* to enable automatic Administrator approval.
4. Click *OK* to save the changes.

10.7 Customizing the Cost Update Schedule

Roles that Can Perform This Task: Cloud Administrator

When a business service is deployed, Cloud Manager takes a snapshot of the business service's workloads. This snapshot maintains a record of the service levels, resource allocations, and costs associated with the workloads at deployment time.

For business services with an expiration date, the snapshot is updated only if the business service is changed and redeployed. At that time, Cloud Manager takes a new snapshot that reflects any changes to the service levels, resource allocations, and costs associated with the workloads. For example, if the cost of the resources has increased since the first deployment, the second deployment reflects the cost increase.

For business services without an expiration date, you can determine when workload snapshots are updated so that non-expiring business services accurately reflect changes in the service levels, resource allocations, and costs associated with their workloads. You can update workloads either weekly or monthly. The update occurs on the last day of the selected period. You can also force an immediate update.

1. On the main navigation bar, click *Configuration*, then click *Cost Updates*.
2. Select *Weekly* or *Monthly* for the update frequency.
3. If you want to force an update immediately, click *Update Now.*
4. Click *OK*.

10.8 Customizing the Cloud Manager Console Interface

Roles that Can Perform This Task: Cloud Administrator

The Cloud Manager console has two interface settings that you can customize: *Currency* and *Workload Dialog*. *Currency* determines the display symbol that is used in Cloud Manager currency fields. *Workload Dialog* determines which tabs (*Windows Settings*, *Windows Licensing*, *Linux Settings*, and *Networks*) are displayed when creating and managing business service workloads.

1. On the main navigation bar, click *Configuration*, then click *User Interface*.
2. Under *Currency*, select the currency symbol you want to use.

   This setting affects only the display symbol. It does not affect the format of the currency fields. All fields are formatted as 00.00 regardless of the currency symbol. In addition, changing the currency symbol does not perform any currency conversion on existing costs (workloads, workload templates, resources, and so on). For example, if you change from United States Dollar (USD) to Euro (EUR), $50.00 simply becomes €50.00.

   Any users who log in after the change see the new currency symbol. For you to see the change, you must log out and then log in again.
3. Under *Workload Dialog*, select the workload tabs you want hidden.

   A workload contains *Windows Settings, Windows Licensing*, *Linux Settings*, and *Networks* tabs that must be configured when requesting a business service. By default, these are displayed so that the user can fill in the required information.

   If you don’t want to require the user to provide this information, you can hide any or all of the tabs. The Build Administrator or a Cloud Administrator must then provide the information when completing the pre-build configuration task for the requested business service.

   All organizations inherit this setting; however, you can override the inherited setting at each organization.
4. Under *Workload Dialog*, specify whether the entering the virtual machine name is to be mandatory for pre-configuration.
5. Click *OK* to save your change.

10.9 Setting Up the Web Application Launcher (Launch Pad)

Roles that Can Perform This Task: Cloud Administrator

Cloud Manager makes it easy for you to add links in its Web Console that point to external Web applications or internal corporate pages. Some customers use the widget to make data center management applications (such as firewall configuration, incident reporting, network monitoring, and so on) more readily available.

The default name for this widget is *Launch Pad*, although you can rename it as you like for display in the interface. Some of the features of the widget include:

- Items you add to Launch Pad are displayed without the need to close and reopen Cloud Manager.
- There is no limit to the number or size of Launch Pad items you can define: the drop-down auto-scales accordingly.
- You can define Launch Pad items for use by the individuals who log on to Cloud Manager: system users, organization users, or both.
- Launch pad items can be embedded for launch inside the Cloud Manager user interface or launched in an external browser window. The Launch Pad window size is configurable.
- You can use hover tooltips for each Launch Pad item to provide information to end users about that item.

The figure below shows a sample Launch Pad configured with items that, when selected, launch a window that exposes either an application or a Web page:



10.9.1 Configuring the Widget

When you launch the Cloud Manager Application Server Console the first time following installation, Cloud Manager creates the Launch Pad configuration file at `<NCM-writable_location>/cloudmanager/plugins/cloudManager/launchPad/launchpad.xml`.

NOTE: This path name is case-sensitive.

By configuring this file, you can control whether the Launch Pad is visible in the interface, what it is to be called, the Web links and images that appear in the interface, and the type, size, and position of the window that is displayed for each item you configure for the interface. If you choose not configure the file, the widget is not implemented in the console.

A sample of the configurable section of the file is included below:

```
<LaunchPad>
   <LaunchPadTitle></LaunchPadTitle>
   <LaunchPadItems>
      <LaunchPadItem>
         <profile></profile>
         <title></title>
         <description></description>
         <link></link>
         <image></image>
         <image_url></image_url>
         <show_external></show_external>
         <size></size>
         <position></position>
         <show_scrollbars></show_scrollbars>
      </LaunchPadItem>
   </LaunchPadItems>
</LaunchPad>
```

The following table provides a description of each of the configurable properties.

Property	Description or Effect
`<LaunchPadTitle>`	If you provide no value for this property, the widget displays `Launch Pad` as the default access link on the console toolbar. This default name is fully localized.
`<profile>`	You can define Launch Pad items to display for organization users (`user`), as system (`admin`) users, or both (`admin`, `user`).
`<title>`	The title of a `LaunchPadItem` displayed below its icon. There is no character limit.
`<description>`	As the mouse hovers over an item, this description displays in “tooltip” format.
`<link>`	Provide an accessible URL as a hyperlink to the application or Web page.
`<image>`	Provide the URL to the image or icon file that you want to use as a representation of this item. You can use any size image; Cloud Manager scales the image in the widget to a predefined size. We recommend you use an image with width and height approximately the same – a minimum of 45 pixels in each dimension. You can also provide an image the using `<image_url>` property, discussed below. f you specify both properties, Cloud Manager defaults to `<image_url>`.
`<image_url>`	Provide the path to image or icon file that you want to use as a representation of this item. You can use any size image; Cloud Manager scales the image in the widget to a predefined size. Use the `<size>` property if you want to define the size yourself. We recommend you use an image with width and height approximately the same – a minimum of 45 pixels in each dimension. You can also provide an image the using `<image>` property, discussed above. f you specify both properties, Cloud Manager defaults to `<image_url>`.

Property	Description or Effect
<show_external>	Use this property to determine the type of window opened when you click the image for the Launch Pad item: ◆ Specifying the false value for this property displays a popup-like, “Cloud Manager-embedded” window that can be sized, moved, and maximized as needed. The window “floats” over the console interface, even on top of new Cloud Manager dialogs that open after it. Embedded windows do not prevent access to other Cloud Manager user interface items. In conjunction with <size> and <position>, you could use this type of window as a kind of monitor for business services. Embedded windows are not persistent. When you close the Cloud Manager console, the embedded window is also closed. ◆ Specifying the true value for this property displays an “external” window that is actually another browser instance and is less integrated with Cloud Manager. The Cloud Manager Help system uses this type of window. This window can be sized and moved as needed. If you do not specify a value for this property, this external window is the default. If the URL of a Launch Pad item points to a secure Web page, that page may issue a certificate challenge to the browser. While an external window allows the user to accept the certificate, embedded windows do not. Make sure you point embedded Launch Pad items only to pages that do not issue a certificate challenge.
<size>	The parameters of the window size property are <size>*width, height*</size>: ◆ For embedded windows, the dimensions can be specified in both pixels and percents. Examples: 200 or 200px or 20%. ◆ For external windows, the dimensions must be specified in pixels. Examples: 200 or 200px. ◆ The parameters can be separated by a comma or by a space ◆ HTML browsers restrict the minimum size of a window. Every browser has different size restrictions. For example, Internet Explorer limits windows to a minimum of 200 x 200 pixels.
<position>	The parameters of the window position property are relative to a static definition of the window placement: <position>*left, top*</size>. ◆ For embedded windows, the values are relative to the top-left of the Cloud Manager Web UI (below the browser’s toolbars). You can specify the values as pixels or percents. Examples:100 or 100px or 80%. **NOTE:** For positions specified using percentage, be careful not to position an item off the edge of the screen For external windows, the values are relative to the top-left of the Cloud Manager browser window (above all tool and menu bars)and are absolute pixel coordinates relative to that window. Examples: 100 or 100px. ◆ The parameters can be separated by a comma or by a space ◆ If you do not provide values for this property, the window is centered in the display.

Property	Description or Effect
`<show_scrollbars>`	You can specify a `true` or a `false` value for this property. True is the default. Activating this property more fully integrates Launch Pad items into the Cloud Manager Web UI: an item properly sized and positioned, with scroll bars removed, appears to be part of the Cloud Manager interface. **NOTE:** Not all scrollbars can be removed. When an embedded window contains content from a different origin (that is, a server domain other than the Cloud Manager server's domain), Cloud Manager is not granted access to hide the scroll bars. If the contents force a larger width than that specified for the embedded window, scroll bars still appear.

If an exception occurs while parsing `launchpad.xml`, the Launch Pad widget is not activated. You can examine the Cloud Manager Application Server log to determine if there are problems with reading the file. We recommend you load the file into an XML validator, such as the Firefox browser (or many other online validators), to ensure that the file contents are properly formed.

11 Setting Up and Managing Users

The following sections provide information to help you manage the users and user groups in your Cloud environment:

- Section 11.1, “Creating User Accounts,” on page 101
- Section 11.2, “Enabling Automatic User Creation,” on page 105
- Section 11.3, “Creating User Groups,” on page 107
- Section 11.4, “Assigning Roles to Users and Groups,” on page 111
- Section 11.5, “Deleting Users,” on page 112
- Section 11.6, “Deleting User Groups,” on page 112

11.1 Creating User Accounts

Access to Cloud Manager requires a Cloud Manager user account. Through the account, a user receives rights to perform various roles in the Cloud Manager system, in an organization, or in both.

There are two types of user accounts: System and Organization. A System account enables a user to be assigned system-level roles (Approver, Build Administrator, Catalog Manager, Cloud Administrator, and Zone Administrator) and organization-level roles (Business Group Viewer, Business Service Owner, Organization Manager, and Sponsor). You can also create accounts for Organization users. Organization users can be assigned organization-level roles only.

You can create users by manually entering information or by importing information from your LDAP authentication source.

- Section 11.1.1, “Manually Creating System and Organization Users,” on page 101
- Section 11.1.2, “Manually Importing System Users from LDAP,” on page 102
- Section 11.1.3, “Manually Importing Organization Users from LDAP,” on page 103

11.1.1 Manually Creating System and Organization Users

Roles that Can Perform This Task: Cloud Administrator, Organization Manager (Organization users only)

The following steps explain how to create users by manually entering their information. For information about creating users by importing their information from your LDAP authentication source, see “Manually Importing System Users from LDAP” on page 102 and Section 11.1.3, “Manually Importing Organization Users from LDAP,” on page 103.

1. On the main navigation bar, click *Users*.
2. Click the *Users* tab, then click *Create* to display the Create User dialog box.
3. Provide the following details to define the user:

   Full Name: Specify the user’s full name as you want it to appear in Cloud Manager.

E-Mail Address: Specify the user's email address as defined in their LDAP authentication account. If necessary, you can specify more than one address; use commas to separate addresses.

The e-mail address enables the Cloud Manager system to send messages (tasks, notifications, and so forth) to the user as needed.

If LDAP is being used for authentication (without Access Manager or Cloud Security Services), the e-mail address is also used for login.

Phone Number: This field is optional. Specify a contact number if desired.

4 Select the user's scope:

Organization: An organization scope enables the user to perform roles within a specific organization. The roles are Business Group Viewer, Business Service Owner, Organization Manager, and Sponsor.

To give the user an organization scope, select *Organization*, then select the organization in which to place the user.

System: A system scope enables the user to administer the Cloud Manager system. The roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and Zone Administrator. In addition, a System user can be given any of the organization roles.

5 Determine cost visibility for the user:

If you want business service owners to see the costs associated with their workload templates, select *Always show costs*. If this check box is not selected, and the user's visibility in the organization and business group is not set, all costs are hidden for the user.

6 Add the user to user groups.

When you add a user to a group, the user inherits the roles assigned to the group.

6a Click the *Membership* tab.

6b Click *Add*, select the desired user groups, then click *OK*.

You can Shift-click and Ctrl-click to select multiple groups.

7 Click *Save* to add the user to the *Users* list.

8 To assign roles to the user, see Assigning Roles to Users and Groups.

11.1.2 Manually Importing System Users from LDAP

Roles that Can Perform This Task: Cloud Administrator

The following steps explain how to create System users by importing information from your LDAP authentication source. For information about creating System users by manually entering information, see "Manually Creating System and Organization Users" on page 101.

1 On the main navigation bar, click *Configuration*.

2 Click *System Users*, click *Members*, then click *Import*.

3 Authenticate to the LDAP directory:

3a Click the *LDAP* tab.

3b In the *LDAP Location* section, fill in the following fields:

Host: Specify the FQDN (fully qualified domain name) or IP address of the host machine running the LDAP server. For example, `ldap.mycompany.com` or `123.45.67.8`.

Port: Specify the TCP port (on the host machine) where the LDAP server is listening for LDAP connections. The standard port for non-SSL connections is 389. The standard port for SSL connections is 636.

Use SSL: If the Cloud Manager Application Server is configured for an SSL connection to the LDAP server, select this option to enable the secure connection.

3c In the *Search Bind Account* section, fill in the following fields:

DN: Specify the distinguished name of an account that has search rights to the directory location from which you want to import users. For example, `cn=Administrator,cn=Users,dc=MyCompany,dc=com`

Password: Specify the password for the account.

Confirm Password: Confirm the password for the account.

3d Click *Test Connection*.

If the connection is successful, the Test Status is displayed as *Passed*. If the connection is not successful, validate the connection information and try again.

4 Import users:

4a Click the *Import* tab.

4b Click *Add*.

When you click *Add*, a new import entry is added to the list. You use the fields below the list to define the entry.

4c In the *DN* field, use standard LDAP notation (`ou=provo,dc=netiq,dc=com`) to specify the distinguished name for the target container or object, then click *Validate*.

If you specify a container, all users located within the container are imported. If you only want to import one user, specify the DN of the user object.

4d If you specified a container for import, select *Users*.

4e If you specified a container for import, select *Scan Tree* if you want to import users located in its subcontainers.

4f Click *Import*.

The imported users are added to the *Members* list. Users are identified by the  icon.

5 Click *OK* to close the System Configuration dialog box.

6 To assign roles to a user, see Assigning Roles to Users and Groups.

11.1.3 Manually Importing Organization Users from LDAP

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

The following steps explain how to create Organization users by importing information from your LDAP authentication source. For information about creating Organization users by manually entering information, see “Manually Creating System and Organization Users” on page 101.

1 On the main navigation bar, click  *Organizations*.

2 Click the *Organizations* tab, select the target organization for the import, click *Edit* to display the Edit Organization dialog box.

3 On the *Users* tab, click *Members*, then click *Import*.

4 Authenticate to the LDAP directory:

4a Click the *LDAP* tab.

4b In the *LDAP Location* section, fill in the following fields:

Host: Specify the FQDN (fully qualified domain name) or IP address of the host machine running the LDAP server. For example, `ldap.mycompany.com` or `123.45.67.8`.

Port: Specify the TCP port (on the host machine) where the LDAP server is listening for LDAP connections. The standard port for non-SSL connections is 389. The standard port for SSL connections is 636.

Use SSL: If the Cloud Manager Application Server is configured for an SSL connection to the LDAP server, select this option to enable the secure connection.

4c In the *Search Bind Account* section, fill in the following fields:

DN: Specify the distinguished name of an account that has search rights to the directory location from which you want to import users. For example, `cn=Administrator,cn=Users,dc=MyCompany,dc=com`

Password: Specify the password for the account.

Confirm Password: Confirm the password for the account.

4d Click *Test Connection*.

If the connection is successful, the Test Status is displayed as *Passed*. If the connection is not successful, validate the connection information and try again.

5 Import users:

5a Click the *Import* tab.

5b Click *Add*.

When you click *Add*, a new import entry is added to the list. You use the fields below the list to define the entry.

5c In the *DN* field, use standard LDAP notation (`ou=provo,dc=netiq,dc=com`) to specify the distinguished name for the target container or object, then click *Validate*.

If you specify a container, all users located within the container are imported. If you only want to import one user, specify the DN of the user object.

5d If you specified a container for import, select *Users*.

5e If you specified a container for import, select *Scan Tree* if you want to import users located in its subcontainers.

5f Click *Import*.

The imported users are added to the *Members* list. Users are identified by the  icon.

6 Assign roles to a user.

An Organization user can be assigned roles at the organization level, business group level, or business service level. If you want to assign an imported user a role at the organization level, continue with the following steps. If you want to assign roles at the other two levels, exit the dialog box and see Assigning Roles to Users and Groups.

Users must be given roles in order to do anything in the organization. There are six roles that apply at the organization level: Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, and Sponsor.

Role assignments at the organization level are inherited by the organization's business groups. For example, if you give a user the Business Service Owner role for an organization, the user can create business services for any business group in the organization. If you want to limit the user to a role in specific business group, you must make the role assignment in the business group.

6a Click the role (*Approver*, *Build Administrator*, *Business Group Viewer*, *Business Service Owner, Organization Manager*, or *Sponsor*) that you want to assign to a user.

6b Click *Add*.

Depending on the role that you are adding, the selection dialog box can contain two lists: *Members* and *System Users*. The *Members* list includes all members of the organization and the *System Users* list includes all Cloud Manager System users.

6c Select the users you want to add, then click *OK*.

You can Shift-click and Ctrl-click to select multiple users.

7 Click *Save* to close the Edit Organization dialog box.

11.2 Enabling Automatic User Creation

You can enable user accounts to be created automatically the first time valid LDAP users log in to the Cloud Manager Application Console. This process requires you to associate users' domain names with the Cloud Manager system (for importing System users) or with organizations (for importing Organization users). For example, if you associated the `netiq.com` domain name with your system, any user who logged in with `netiq.com` in their e-mail address would be made a System user.

The following sections explain how to set up automatic user creation and how to automate role assignments for automatically imported users:

- Section 11.2.1, "Setting Up Automatic User Creation for System Users," on page 105
- Section 11.2.2, "Setting Up Automatic User Creation for Organization Users," on page 106
- Section 11.2.3, "Automating Role Assignments to Automatically Imported Users," on page 106

11.2.1 Setting Up Automatic User Creation for System Users

Roles that Can Perform This Task: Cloud Administrator

1 On the main navigation bar, click *Configuration*.

2 Click *System Users*.

3 In the *Domains* field, specify the email domains that you to be imported into to the system.

For example, if you want all users who log in with email addresses that include the `netiq.com` or `novell.com` domain names, specify `netiq.com,novell.com`. Use a comma to separate domain names.

4 Click *OK* to save your changes.

11.2.2 Setting Up Automatic User Creation for Organization Users

Roles that Can Perform This Task: Cloud Administrator

1. On the main navigation bar, click *Organizations*.
2. On the *Organizations* tab, select the organization for which you want to import users, then click *Edit*.
3. In the *Domains* field, specify the email domains that you want imported into the organization.

   For example, if you want all users who log in with email addresses that include the `suse.com` domain name, specify `suse.com`. If you specify multiple domains, use a comma to separate domain names.
4. Click *OK* to save your changes.

11.2.3 Automating Role Assignments to Automatically Imported Users

Roles that Can Perform This Task: Cloud Administrator

When a user is automatically imported, his or her user account is created without any role assignments. You can manually assign roles to the user after the account is created, but this negates much of the administrative benefit gained by enabling the automatic import.

To receive the maximum benefit of automatic user creation, you can assign roles to users through the use of LDAP user groups. By assigning roles to LDAP user groups, you can ensure that LDAP users who are members of those groups automatically inherit those roles when they are imported.

To automate role assignments for automatically imported users:

1. In your LDAP source, create the LDAP user groups you want.

   For example, in the LDAP directory used for authenticating System users, you could create an LDAP user group for Cloud Administrators, another for Zone Administrators, and another for Build Administrators. In the LDAP directory used for authenticating an organization's users, you could create LDAP user groups for Organization Managers and Business Service Owners.
2. Add the appropriate LDAP users to each LDAP user group.

   For example, if you created a Business Service Owners group, add the users who are Business Service Owners for the organization.
3. Add the LDAP user groups to Cloud Manager using one of the following methods:
   - Import the user group information from LDAP. For instructions, see "Importing System User Groups from LDAP" on page 108 and "Importing Organization User Groups from LDAP" on page 109.
   - Create the user groups by manually adding group information, including the distinguished name of the user group in LDAP. For instructions, see "Manually Creating System and Organization User Groups" on page 107.
4. Assign roles to the user groups. For instructions, see "Assigning Roles to Users and Groups" on page 111.

11.3 Creating User Groups

Rather than assign roles to individual users, you can create user groups and assign roles to the user groups. Users who are added to a group inherit the group's roles.

User group roles are cumulative. If you add a user to a group, the user retains its directly assigned roles and also gains the roles inherited from the group.

As with users, there are two types of user groups: System and Organization. A System group can be assigned system-level roles (Approver, Build Administrator, Catalog Manager, Cloud Administrator, and Zone Administrator) and organization-level roles (Business Group Viewer, Business Service Owner, Organization Manager, and Sponsor). An Organization user group can be assigned organization-level roles only.

You can create user groups by manually entering information or by importing information from your LDAP authentication source.

- Section 11.3.1, "Manually Creating System and Organization User Groups," on page 107
- Section 11.3.2, "Importing System User Groups from LDAP," on page 108
- Section 11.3.3, "Importing Organization User Groups from LDAP," on page 109

11.3.1 Manually Creating System and Organization User Groups

Roles that Can Perform This Task: Cloud Administrator, Organization Manager (Organization users only)

The following steps explain how to create user groups by manually entering information. For information about creating user groups by importing information from your LDAP authentication source, see "Importing System User Groups from LDAP" on page 108 and Section 11.1.3, "Manually Importing Organization Users from LDAP," on page 103.

1. On the main navigation bar, click *Users*.
2. Click the *User Groups* tab, then click *Create* to display the Create User Group dialog box.
3. Provide the following details to define the user group:

   Full Name: Specify the group's full name as you want it to appear in Cloud Manager.

   E-Mail Address: This field is optional. If you enter an email address, any messages generated for the group's roles are sent to the email address. If you don't enter an email address, the messages are sent to the group members' addresses.
4. In the *Scope* field, select *System*.
5. In the *Type* field, select the group's type:
   - **LDAP DN:** Select this option to specify an LDAP group. The group's membership is maintained in the LDAP source. You cannot add users to the group in Cloud Manager.

     Use standard LDAP notation to specify the distinguished name of the user group in the LDAP source (for example, `cn=orgmanagers,dc=provo,dc=netiq,dc=com`).
   - **Cloud Manager:** Select this option to create a user group that exists only in Cloud Manager. You maintain the group membership in Cloud Manager. The group can include both users and other groups (including LDAP user groups).

6 Add members to the group:

6a Click the *Membership* tab.

6b Click *Members*, then click *Add* to display the Add Members dialog box.

6c Select the users and user groups you want to add to the group.

You can Shift-click and Ctrl-click to select multiple users and groups.

6d Click *OK* to add the users and user groups to the Members list.

7 Click *Save*.

8 To assign roles to the user, see Assigning Roles to Users and Groups.

11.3.2 Importing System User Groups from LDAP

Roles that Can Perform This Task: Cloud Administrator

The following steps explain how to create System user groups by importing information from your LDAP authentication source. For information about creating Organization user groups by manually entering information, see "Manually Creating System and Organization User Groups" on page 107.

An imported user group's membership is maintained in the LDAP authentication source. Any users who are members of the user group in the LDAP source receive the roles that are assigned to the user group in Cloud Manager.

An LDAP user group's members are not imported to Cloud Manager and do not display in the group's *Members* list. In addition, you cannot manually add users or user groups to an imported group.

1 On the main navigation bar, click *Configuration*.

2 Click *System Users*, click the *Members* tab, then click *Import*.

3 Authenticate to the LDAP directory:

3a In the Import from Directory dialog box, click the *LDAP* tab.

3b In the *LDAP Location* section, fill in the following fields:

Host: Specify the FQDN (fully qualified domain name) or IP address of the host machine running the LDAP server. For example, `ldap.mycompany.com` or `123.45.67.8`.

Port: Specify the TCP port (on the host machine) where the LDAP server is listening for LDAP connections. The standard port for non-SSL connections is 389. The standard port for SSL connections is 636.

Use SSL: If the Cloud Manager Application Server is configured for an SSL connection to the LDAP server, select this option to enable the secure connection.

3c In the *Search Bind Account* section, fill in the following fields:

User DN: Specify an account that has read rights to the directory location from which you want to import users. For example, `cn=Administrator,cn=Users,dc=MyCompany,dc=com`

Password: Specify the password for the account.

Password Confirm: Confirm the password for the account.

3d Click *Test Connection*.

If the connection is successful, the Test Status is displayed as *Passed*. If the connection is not successful, validate the connection information and try again.

4 Import user groups:

4a Click the *Import* tab.

4b Click *Add*.

When you click *Add*, a new import entry is added to the list. You use the fields below the list to define the entry.

4c In the DN field, use standard LDAP notation (`ou=provo,dc=netiq,dc=com`) to specify the distinguished name for the target container or object, then click *Validate*.

If you specify a container, all user groups located within the container are imported. If you only want to import one user group, specify the DN of the user group object.

4d If you specified a container for import, select *Groups*.

4e If you specified a container for import, select *Scan Tree* if you want to import user groups located in its subcontainers.

4f Click *Import*.

The imported user groups are added to the *Members* list. User groups are identified by the  icon.

5 Click *Save* to close the System Configuration dialog box.

6 To assign roles to a user group, see Assigning Roles to Users and Groups.

11.3.3 Importing Organization User Groups from LDAP

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

The following steps explain how to create Organization user groups by importing information from your LDAP authentication source. For information about creating Organization user groups by manually entering information, see "Manually Creating System and Organization Users" on page 101.

An imported user group's membership is maintained in the LDAP authentication source. Any users who are members of the user group in the LDAP source receive the roles that are assigned to the user group in Cloud Manager.

An LDAP user group's members are not imported to Cloud Manager and do not display in the group's *Members* list. In addition, you cannot manually add users or user groups to an imported group.

1 On the main navigation bar, click  *Organizations*.

2 Click the *Organizations* tab, select the target organization for the import, click *Edit* to display the Edit Organization dialog box.

3 On the *Users* tab, click *Members*, then click *Import*.

4 Authenticate to the LDAP directory:

4a In the Import from Directory dialog box, click the *LDAP* tab.

4b In the *LDAP Location* section, fill in the following fields:

Host: Specify the FQDN (fully qualified domain name) or IP address of the host machine running the LDAP server. For example, `ldap.mycompany.com` or `123.45.67.8`.

Port: Specify the TCP port (on the host machine) where the LDAP server is listening for LDAP connections. The standard port for non-SSL connections is 389. The standard port for SSL connections is 636.

Use SSL: If the Cloud Manager Application Server is configured for an SSL connection to the LDAP server, select this option to enable the secure connection.

4c In the *Search Bind Account* section, fill in the following fields:

DN: Specify the distinguished name of an account that has search rights to the directory location from which you want to import users. For example, `cn=Administrator,cn=Users,dc=MyCompany,dc=com`

Password: Specify the password for the account.

Confirm Password: Confirm the password for the account.

4d Click *Test Connection*.

If the connection is successful, the Test Status is displayed as *Passed*. If the connection is not successful, validate the connection information and try again.

5 Import user groups:

5a Click the *Import* tab.

5b Click *Add*.

When you click *Add*, an new import entry is added to the list. You use the fields below the list to define the entry.

5c In the *DN* field, use standard LDAP notation (`ou=provo,dc=netiq,dc=com`) to specify the distinguished name for the target container or object, then click *Validate*.

If you specify a container, all user groups located within the container are imported. If you only want to import one user group, specify the DN of the user group object.

5d If you specified a container for import, select *Groups*.

5e If you specified a container for import, select *Scan Tree* if you want to import users located in its subcontainers.

5f Click *Import*.

The imported user groups are added to the *Members* list. User groups are identified by the  icon.

6 Assign roles to a user group.

An Organization user group can be assigned roles at the organization level, business group level, or business service level. If you want to assign an imported user group a role at the organization level, continue with the following steps. If you want to assign roles at the other two levels, exit the dialog box and see Assigning Roles to Users and Groups.

User groups must be given roles in order for group members to do anything in the organization. There are six roles that apply at the organization level: Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, and Sponsor.

Role assignments at the organization level are inherited by the organization's business groups. For example, if you give a group the Business Service Owner role for an organization, the group members can create business services for any business group in the organization. If you want to limit the user group to a role in specific business group, you must make the role assignment in the business group.

6a Click the role (*Approver*, *Build Administrator*, *Business Group Viewer*, *Business Service Owner, Organization Manager*, or *Sponsor*) that you want to assign to a user group.

6b Click *Add*.

Depending on the role that you are adding, the selection dialog box can contain two lists: *Members* and *System Users*. The *Members* list includes all members of the organization and the *System Users* list includes all Cloud Manager System users.

6c Select the user groups you want to add, then click *OK*.

You can Shift-click and Ctrl-click to select multiple groups.

7 Click *Save* to close the Edit Organization dialog box.

11.4 Assigning Roles to Users and Groups

Roles that Can Perform This Task: Cloud Administrator, Organization Manager (Organization users only)

System users and user groups can be assigned both system-level roles (Approver, Build Administrator, Catalog Manager, Cloud Administrator, System View, and Zone Administrator) and organization-level roles (Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor). Organization users can be assigned organization-level roles only.

For role descriptions, see "Cloud Manager Roles" in the *NetIQ Cloud Manager 2.3 Product Overview*.

1 On the main navigation bar, click *Users*.

2 To assign a role to a user, click the *Users* tab, select the user, then click *Edit* to display the Edit User dialog box.

or

To assign a role to a user group, click the *User Groups* tab, select the user group, then click *Edit* to display the Edit User Group dialog box.

3 (System user or group only) Assign system-level roles.

The system-level roles are Approver, Build Administrator, Catalog Manager, Cloud Administrator, and System View. These roles can be assigned only to System users or groups.

NOTE: System scope also implies the Zone Administrator role, though it is not explicitly listed. Instead, specific zones are associated to these users or user groups, making the Zone Administrator role implicit.

3a To assign the Approver, Build Administrator, Catalog Manager, or Cloud Administrator role, click the *System* tab, click *Add*, select the desired roles, then click *OK*.

3b To assign the Zone Administrator role, click the *Zone* tab, click *Add*, select the desired zone, then click *OK*.

4 Assign organization-level roles.

The organization-level roles are Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, Sales Manager, and Sponsor. The Approver and Build Administrator roles can be assigned only to System users and groups. The Sales Manager role can be assigned only to Organization users and groups. The other roles can be assigned to both System and Organization users and groups.

Several of the roles can be assigned at the organization, business group, or business service level. For example, you can make a user a Sponsor for a business group, in which case the user can approve requests for business services from that business group only. Or, you can make the user a Sponsor for the organization, in which case the user can approve requests for all business services in the organization.

4a Click the *Organization* tab to add a role at the organization level, click the *Business Group* tab to add a role at the business group level, or click the *Business Service* tab to add a role at the business service level.

4b Click the role that you want to assign.

For example, if you selected the *Business Group* tab and you want to enable the user or group to create business services for the business group, click *Business Service Owner*.

4c Click *Add*, select the object (organization, business group, or business service) to which you want the role to apply, then click *OK* to add it to the list.

5 When you have finished assigning roles, click *Save* to save the role changes.

11.5 Deleting Users

Roles that Can Perform This Task: Cloud Administrator, Organization Manager (Organization users only)

When you delete a user, be aware of the following:

- The user's deployed business services remain deployed. Other users (Cloud Administrators, Business Service Owners, and so forth) who had access to the business services continue to have access. If the user was the sole owner of a business service, you can assign other users as owners of the business service, either before or after the user is deleted.
- The user's business service requests remain in progress. As with deployed business services, other users (Cloud Administrators, Business Service Owners, and so forth) who had access to the requests continue to have access. If the user was the sole owner of a business service request, you can assign other users as owners of the business service, either before or after the user is deleted.
- The user's claimed tasks must be released (unclaimed) or claimed by another user, either a Cloud Administrator or another user who has the same role as the deleted user. The tasks can be claimed before or after the user is deleted.

To delete a user:

1. On the main navigation bar, click *Users*.
2. Click the *Users* tab, select the user to delete, then click *Delete*.
3. Click *Yes* to confirm the deletion.

11.6 Deleting User Groups

Roles that Can Perform This Task: Cloud Administrator, Organization Manager (Organization user groups only)

When you delete a user group, the group members are not deleted, but they lose any roles that they inherited through membership in the group.

1. On the main navigation bar, click *Users*.
2. Click the *User Groups* tab, select the user group to delete, then click *Delete*.
3. Click *Yes* to confirm the deletion.

12 Managing the Workload Template Catalog

The following sections provide information to help you manage the catalog of workload templates in your Cloud environment:

- Section 12.1, “Creating Workload Templates,” on page 113
- Section 12.2, “Assigning Workload Templates to Organizations and Business Groups,” on page 116
- Section 12.3, “Modifying Workload Templates,” on page 118
- Section 12.4, “Deleting Workload Templates,” on page 118
- Section 12.5, “Add-on Applications,” on page 119
- Section 12.6, “Add-On Services,” on page 121
- Section 12.7, “Legal Agreements,” on page 124
- Section 12.8, “Resource Pricing Packages,” on page 127

12.1 Creating Workload Templates

Business service workloads are created from workload templates. Each workload template identifies a VM template and the customizations (such as increased CPUs or decreased memory) that you want applied when creating a workload from the template.

The VM templates you can use come from your Cloud Manager zones. A single VM template can be used in multiple workload templates.

1. On the main navigation bar, click *Getting Started*, then click *Create Workload Templates* (in the *Set Up Your Cloud Environment* list).

   or

   On the main navigation bar, click *Catalog,* click the *Workload Templates* tab, then click *Create*.
2. If your Cloud Manager system has multiple zones, the Select Zone dialog box is displayed. Select the zone that contains the VM template you want, then click *OK* to display the Create Workload Template dialog box.
3. In the *General* section, provide the following information for the workload template:

   Name: Specify a unique name for the workload template. Users see this name (along with the description, zone, and operating system) when selecting the workload template they use to create their business service workload.

   You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Setup Cost: Specify the cost associated with setting up the workload. This is a one-time fee.

   License Cost: Specify the license costs associated with the workload’s software. This cost is per month.

Zone: Displays the zone you selected for the workload template. You cannot change this setting.

Creation Date: Displays the current date. You cannot change this setting.

Description: Provide any additional information to identify the workload template. Users see this description (along with the name, zone, and operating system) when selecting the workload template from which to create business service workload.

4 In the *Virtual Machine Settings* section:

4a In the *VM Template* list, select a VM template.

After you select a VM template, the template's operating system and hypervisor information is displayed. You cannot change these settings.

The VM template's resource information (CPUs, memory, network interface cards, and disks) is also displayed. You can customize this information as necessary.

4b Customize the settings to increase or decrease the workload resources:

Number of CPUs: Select the number of CPUs for the workload. Some hypervisor technologies do not support more than 8 CPUs per workload, so the maximum number allowed is 8.

Memory: Select the megabytes (MB) of RAM to allocate to the workload.

Number of NICs: Select the number of network interface cards (NICs) to allocate to the workload.

Disks Summary: Displays the size of the mandatory workload disks and the number of optional workload disks defined in the template, Mandatory workload disks are always created. They are inherited from the VM template or manually defined on the *Disks* tab. Optional workload disks are available to the user but are created only if the user specifies sizes for the disks.

You add disks on the Disks page (see Step 5 below) and the disk information is then displayed in these fields.

4c By default, each resource setting is unlocked, which means that users can change it when creating workloads from the template. If you want to prevent users from changing a setting, select the  check box.

5 If you want to add disks to the workload template, or if you want to specify whether or not the current disks should be included when calculating the cost of the workload:

5a Click the *Disks* tab.

The template can include up to 10 additional disks.

5b Click *Add* to add a disk to the list.

5c Make sure the disk is selected in the list, then specify a size (in the *Size* field below the list) to make the disk mandatory.

or

Leave the size set to 0 to make the disk optional.

The maximum size per disk is 1024 GB (1 TB). If you specify the size when adding a disk, the disk becomes mandatory. Users cannot remove or change a mandatory disk. If you add a disk with the size set to 0, the disk becomes optional. Users can leave the size at 0, in which case the disk is not created with the workload, or they can specify a size and create the disk.

5d In the *Cost* field, select the check box if you want to include the cost of the disk in the workload's total cost.

If you don't enable the *Cost* option, the disk becomes a free disk and is not included when calculating the cost of the workload.

6 If the template is a Windows-based template and you want to pre-populate some of the Windows settings, click *Windows Settings*, then complete the following steps.

You might not want to pre-populate some settings, such as *Computer Name*, if the template will be used to create multiple workloads. Any settings that you do not pre-populate must be filled in when requesting a new business service (by the user) or when performing the pre-build configuration (by an administrator).

6a Configure the *Domain Settings*:

Computer Name: Specify the computer name for the virtual machine.

Domain or Workgroup: Select *Domain* or *Workgroup*, then specify the name of the domain or workgroup to which you want the virtual machine added.

Domain Administrator User ID: This applies only if you are adding the virtual machine to a domain. Specify a domain administrator user ID that can be used to add the virtual machine to the domain specified in the *Domain* field.

6b Modify the *Installation Settings*:

Run Once Commands: Specify any Windows RunOnce commands that you want run during the first log in to the virtual machine. For information about Windows RunOnce commands, see the Microsoft Windows documentation.

7 If the template is a Windows-based template and you want to pre-populate some of the Windows licensing information, click *Windows Licensing* and fill in the fields., then click *OK* to create the workload template and add it to the list.

Windows Product Key: Specify the product key for the workload's Windows operating system.

If you pre-populate this field in the template, the data is masked from users and cannot be copied.

Registered to Name: Specify an individual, department, company or so forth to whom the Windows operating system software is registered.

You might not want to pre-populate some settings, such as *Windows Product Key*, if the template will be used to create multiple workloads and the key does not cover multiple installations. Any settings that you do not pre-populate must be filled in when requesting a new business service (by the user) or when performing the pre-build configuration (by an administrator).

8 If the template is a Linux-based template and you want to pre-populate some of the Linux settings, click *Linux Settings* and fill in the fields, then click *OK* to create the workload template and add it to the list.

Hostname: Specify the host name for the workload

Domain Name: Specify the domain for the workload (for example, netiq.com or provo.netiq.com).

You might not want to pre-populate some settings, such as *Hostname*, if the template will be used to create multiple workloads. Any settings that you do not pre-populate must be filled in when requesting a new business service (by the user) or when performing the pre-build configuration (by an administrator).

9 (Optional) If you have already created organizations and resource groups that support workload templates, you can associate the workload template with the organizations and business groups of your choice, but you must choose an organization whose resource groups support the type of hypervisor (and VM template) compatible with the workload template.

9a Click *Associations*, then select either the *Organizations* tab or the *Business Groups* tab to open a list.

9b Select (that is, click) the organization or business group you want to associate with this workload template, then click *OK*.

or

Select (that is, Ctrl+click) the organizations or business groups you want to associate with this workload template, then click *OK*.

NOTE: The Workload Templates list displays only the workload templates from the organization where your business group resides. You must assign workload templates at the organization level before they become available for selection at the business group level.

12.2 Assigning Workload Templates to Organizations and Business Groups

The workload template catalog is not directly available to organizations. Any workload templates you want available to an organization must be assigned to the organization. After a workload template is assigned to an organization, you can then make it available for use in all or some of the organization's business groups.

- Section 12.2.1, "Assigning Workload Templates to an Organization," on page 116
- Section 12.2.2, "Assigning Workload Templates to a Business Group," on page 116
- Section 12.2.3, "Removing Workload Template Assignments from an Organization," on page 117
- Section 12.2.4, "Removing Workload Template Assignments from a Business Group," on page 117

12.2.1 Assigning Workload Templates to an Organization

Roles that Can Perform This Task: Cloud Administrator

1. On the main navigation bar, click *Organizations*.
2. In the *Organizations* list, select the target organization, then click *Edit*.
3. Click the *Workload Templates* tab.
4. Click *Add* to display the Add Workload Templates dialog box, select the workload templates to be added, then click *OK*.

   You can Shift-click or Ctrl-click to select multiple workload templates.
5. When you have finished adding templates, click *Save* to save the changes to the organization.

12.2.2 Assigning Workload Templates to a Business Group

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

A business group does not automatically inherit the workload templates assigned to its organization. Any of the organization's workload templates that you want available to the business group must be assigned to it.

1. On the main navigation bar, click *Organizations*, then click the *Business Groups* tab.
2. In the *Business Groups* list, select the target business group, then click *Edit*.

3 Click the *Workload Templates* tab.

4 Click *Add* to display the Add Workload Templates dialog box, select the workload templates to be added, then click *OK*.

   You can Shift-click or Ctrl-click to select multiple workload templates.

5 When you have finished adding templates, click *Save* to save the changes to the business group.

12.2.3 Removing Workload Template Assignments from an Organization

Roles that Can Perform This Task: Cloud Administrator

You can remove a workload template from an organization unless it is currently being used to build a workload for one of the organization's business groups. As soon as the workload is built and deployed, you can remove the workload template.

Removing a workload template assignment from an organization also removes any assignments from its business groups.

1 On the main navigation bar, click *Organizations*.

2 In the *Organizations* list, select the organization from which you want to remove the workload template assignment, then click *Edit*.

3 Click the *Workload Templates* tab.

4 Select the workload template to be removed, click *Remove*, then click *Yes* to confirm the template removal.

   You can Shift-click or Ctrl-click to select multiple workload templates.

5 When you have finished removing templates, click *Save* to save the changes to the organization.

12.2.4 Removing Workload Template Assignments from a Business Group

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

You can remove a workload template from a business group unless it is currently being used to build a workload for the business group. As soon as the workload is built and deployed, you can remove the workload template.

1 On the main navigation bar, click *Organizations*.

2 Click the *Business Groups* tab, select the business group from which you want to remove the workload template assignment, then click *Edit*.

3 Click the *Workload Templates* tab.

4 Select the workload template to be removed, click *Remove*, then click *Yes* to confirm the template removal.

You can Shift-click or Ctrl-click to select multiple workload templates.

5 When you have finished removing templates, click *Save* to save the changes to the business group.

12.3 Modifying Workload Templates

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

You can change workload template settings at any time, even if the template has been used to create workloads. The only restriction is that you cannot specify a different VM template if the workload template is in use by a requested or deployed business service.

Changing a workload template has no immediate effect on deployed workloads. However, if a change is requested for a deployed workload, the workload settings are validated against the new workload template settings. This might require the Business Service Owner to change settings that he or she did not plan to change. For example, suppose that you create a workload template that allocates 4 CPUs. A Business Service Owner creates a workload (with 4 CPUs) from the workload template. You then change the workload template's CPU allocation from 4 to 2. After the change, the Business Service Owner requests a change to the workload's number of disks. When creating the change request, the Business Service Owner is also required to change the CPUs from 4 to 2 because 4 CPUs are no longer supported by the new workload template.

1 On the main navigation bar, click *Catalog*.

2 On the *Workload Templates* tab, select the template to modify, then click *Edit*.

3 Make the desired changes to the template, then click *OK* to save the changes.

For a description of the workload template settings, see "Creating Workload Templates" on page 113 or click the icon in the Edit Workload Template dialog box.

12.4 Deleting Workload Templates

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

You can delete a workload template unless it is currently being used to build a workload. As soon as the workload is built and deployed, you can delete the workload template.

Deleting a workload template has no effect on deployed workloads, even if the Business Service Owner of one of the workloads requests a change to it.

1 On the main navigation bar, click *Catalog*.

2 On the *Workload Templates* tab, select the template to delete, then click *Delete*.

3 Confirm the deletion.

12.5 Add-on Applications

You can make licensed or proprietary applications available to organization users who might find that such applications add value to the business services they request. You can create a selection of such applications in the Cloud Manager catalog, and then associate one or more of these applications to a workload template.

You can define the initial setup costs and the ongoing monthly cost for each application. For example, you might define the setup cost of the application at $25 and its monthly operating costs at $10.

An add-on application can be associated to multiple workload templates. For example, two similar workload templates might offer the same application.

Multiple add-on applications can also be associated to a single workload template. The user, when requesting a business service, could select a workload template with the desired application.

This section includes the following information:

- Section 12.5.1, “Creating an Add-on Application,” on page 119
- Section 12.5.2, “Configuring an Add-on Application,” on page 120

12.5.1 Creating an Add-on Application

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

When you associate an add-on application with a workload template, the setup costs and monthly operating costs are applied to any business service workloads deployed to the resource group’s hosts.

For example, you might define the setup cost of the application at $25 and its monthly operating costs at $10, making the initial month’s cost $35 and $25 monthly thereafter.

An add-on application can be associated to multiple workload templates. For example, two similar workload templates might offer the same application.

1. On the main navigation bar, click *Catalog*.
2. On the Add-ons tab, click *Applications*, then click *Create* to open a Create Add-on Application dialog box.
3. In the *General* section of the dialog box, provide the following details for the add-on application:

   Name: Specify a name for the application. The name should be different from any other application name.

   Monthly Cost: Define the ongoing cost (per month) for using the workload templates associated with the application.

   Setup Cost: Specify the one-time fee associated with the licensing, installation, and configuration of the application.

   Creation Date: Displays the current date.

   Description: Provide any additional information to further identify the application.

   Configuration Instructions: Enter text asking the user of the application how the application should be configured. If you enter text here, the business service owners must provide a reply in order to add the application to a workload.
4. Click *Create* to display the Create Add-On Application dialog box

5 Associate the application with workload templates:

5a Under *Associations*, click the Workload Templates tab.

5b Click *Add* to display the Add Workload Templates dialog box

5c Select the workload templates to add.

You can Shift-click and Ctrl-click to select multiple objectives.

5d Click *OK* to add the selected objectives to the *Service Level Objectives* list.

6 Associate the service level with resource groups:

6a Under *Associations*, click the *Resource Groups* tab.

6b Click *Add* to display the Add Resource Groups dialog box.

6c Select the groups to add.

You can Shift-click and Ctrl-click to select multiple workload templates.

6d Click *OK*.

12.5.2 Configuring an Add-on Application

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

You can configure an add-on application to change its setup and monthly costs or to edit the text that asks the user for specific configuration information for the application. Changing a application might impact the cost of business services currently using the application. For more information, see Chapter 15.10, “Displaying or Hiding Business Service Costs,” on page 183.

1 On the main navigation bar, click *Catalog*.

2 On the Add-ons tab, click *Applications*, select an application from the list, then click *Edit* to open a Configure Add-on Application dialog box.

The following table lists the tasks and actions that you might want to perform as you configure an add-on application in Cloud Manager.

Task	Steps
Change the name or description	1. In the *General* section, modify the following details: **Name:** Specify a name for the application. The name should be different from any other application name. **Description:** Provide any additional information to further identify the application.
Change the one-time setup cost or the monthly operating cost	1. In the *Monthly Cost* field, define the cost (per month) to use the application associated with the workload template. 2. In the *Setup Cost* field, define the one-time cost to install and configure this application associated with the workload template.
Change the application configuration instructions	Change the text you provide to the user that instructs him or her to provide configuration instructions to the build administrator. The user cannot add the application to a workload if he or she doesn’t provide this information. If you leave the field blank, the user is not required to provide any information.

Task	Steps
Associate the application to workload templates	1. In the *Workload Templates* list, click *Add* to display the Add Workload Templates dialog box. 2. Select the workload templates to add. You can Shift-click and Ctrl-click to select multiple workload templates. 3. Click *OK*.
Associate the application to legal agreements	1. In the *Legal Agreements* list, click *Add* to display the Add Legal Agreement dialog box. 2. Select the legal agreements to add. You can Shift-click and Ctrl-click to select multiple legal agreements. 3. Click *OK*.
View an add-on application associated to a workload template	1. In the *Workload Templates* list, select a workload template you want to view. 2. When you are finished viewing the workload template, click *Close*.
View a legal agreement associated to a workload template	1. In the *Legal Agreements* list, select a legal agreement you want to view. 2. When you are finished viewing the legal agreement, click *Close*.
Lock the application to a workload templates, making it required	1. In the *Workload Templates* list, select a workload template you want to lock. 2. Click *Lock* to lock the application to the workload template. This makes the application required for this template.
Remove the application from workload templates	1. In the *Workload Templates* list, select the workload templates you want to unassociate from the application. You can Shift-click and Ctrl-click to select multiple templates. 2. Click *Remove*.
Remove the application from legal agreements	1. In the *Legal Agreements* list, select the legal agreements you want to unassociate from the application. You can Shift-click and Ctrl-click to select multiple agreements. 2. Click *Remove*.

12.6 Add-On Services

As a Cloud Administrator or a Catalog Manager, you can make services available to organization users who might find that such services add value to the business services they request. You can create a selection of such services in the Cloud Manager catalog, and then associate one or more of these services to a workload template.

You can define the initial setup costs and the ongoing monthly cost for each service. For example, you might define the setup cost of the service at $25 and its monthly operating costs at $10.

An add-on service can be associated to multiple workload templates. For example, two similar workload templates might offer the same service.

Multiple add-on services can also be associated to a single workload template. The user, when requesting a business service, could select a workload template with the desired service.

- Section 12.6.1, “Creating an Add-on Service,” on page 122
- Section 12.6.2, “Configuring an Add-on Service,” on page 123

12.6.1 Creating an Add-on Service

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

When you associate an add-on service with a workload template, the setup costs and monthly operating costs are applied to any business service workloads deployed to the resource group’s hosts.

For example, you might define the setup cost of the service at $25 and its monthly operating costs at $10, making the initial month’s cost $35 and $25 monthly thereafter.

An add-on service can be associated to multiple workload templates. For example, two similar workload templates might offer the same service.

1. On the main navigation bar, click *Catalog.*
2. On the Add-ons tab, click *Services*, then click *Create* to open a Create Add-on Service dialog box.
3. In the *General* section of the dialog, provide the following details for the add-on service:

   Name: Specify a name for the service. The name should be different from any other service name.

   Monthly Cost: Define the ongoing cost (per month) for using the workload templates associated with the service.

   Setup Cost: Specify the one-time fee associated with the licensing, installation, and configuration of the service.

   Creation Date: Displays the current date.

   Description: Provide any additional information to further identify the service.

   Configuration Instructions: Enter text asking the user of the service how the service should be configured. If you enter text here, the business service owner must provide a reply in order to add the service to a workload.
4. Click *Create* to display the Create Add-On Service dialog box
5. Associate the service with workload templates:
   - **5a** Under *Associations*, click the Workload Templates tab.
   - **5b** Click *Add* to display the Add Workload Templates dialog box
   - **5c** Select the workload templates to add.

     You can Shift-click and Ctrl-click to select multiple objectives.
   - **5d** Click *OK* to add the selected objectives to the *Service Level Objectives* list.
6. Associate the service level with resource groups:
   - **6a** Under *Associations*, click the *Resource Groups* tab.
   - **6b** Click *Add* to display the Add Resource Groups dialog box.
   - **6c** Select the groups to add.

You can Shift-click and Ctrl-click to select multiple workload templates.

6d Click *OK*.

12.6.2 Configuring an Add-on Service

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

You can configure an add-on service to change its setup and monthly costs or to edit the text that asks the user for specific configuration information for the service. Changing a service might impact the cost of business services currently using the service. For more information, see Chapter 15.10, “Displaying or Hiding Business Service Costs,” on page 183.

1. On the main navigation bar, click *Catalog*.
2. On the Add-ons tab, click *Services*, select a service from the list, then click *Edit* to open a Configure Add-on Service dialog box.

The following table lists the tasks and actions that you might want to perform as you configure an add-on service in Cloud Manager.

Task	Steps
Change the name or description	1. In the *General* section, modify the following details: **Name:** Specify a name for the service. The name should be different from any other service name. **Description:** Provide any additional information to further identify the service.
Change the one-time setup cost or the monthly operating cost	1. In the *Monthly Cost* field, define the cost (per month) to use the service associated with the workload template. 2. In the *Setup Cost* field, define the one-time cost to install and configure this service associated with the workload template.
Change the add-on service configuration instructions	Change the text you provide to the user that instructs him or her to provide configuration instructions to the build administrator. The user cannot add the service to a workload if he or she doesn’t provide this information. If you leave the field blank, the user is not required to provide any information.
Associate the service to workload templates	1. In the *Workload Templates* list, click *Add* to display the Add Workload Templates dialog box. 2. Select the workload templates to add. You can Shift-click and Ctrl-click to select multiple workload templates. 3. Click *OK*.
View an add-on service associated to a workload template	1. In the *Workload Templates* list, select a workload template you want to view. 2. When you are finished viewing the workload template, click *Close*.
Lock the service to a workload templates	1. In the *Workload Templates* list, select a workload template you want to lock. 2. Click *Lock* to lock the service to the workload template. This makes the service required for this template.

Task	Steps
Remove the service from workload templates	1. In the *Workload Templates* list, select the workload templates you want to unassociate from the service. You can Shift-click and Ctrl-click to select multiple templates. 2. Click *Remove*.

12.7 Legal Agreements

As a Cloud Administrator or a Catalog Manager, you can create (or copy and paste) legal agreement text that you can associate with add-on applications or services or to workload templates that use such add-ons. When created, these legal agreements are listed in the Cloud Manager catalog. From the catalog, you can edit a legal agreement that has not yet been accepted by a user as part of a workload or you can view it to check its associations. A legal agreement cannot be edited or deleted after it has been accepted by a user.

The Legal Agreements page lists all of the agreements that have been identified in the Cloud Manager system.

- Section 12.7.1, “Creating a Legal Agreement,” on page 124
- Section 12.7.2, “Editing a Legal Agreement,” on page 125

12.7.1 Creating a Legal Agreement

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

As a Cloud Administrator or a Catalog Manager, you can create (or copy and paste) legal agreement text that you can associate with add-on applications or services or to workload templates that use such add-ons. When created, these legal agreements are listed in the Cloud Manager catalog. From the catalog, you can edit a legal agreement that has not yet been accepted by a user as part of a workload or you can view it to check its associations. A legal agreement cannot be edited or deleted after it has been accepted by a user.

An add-on service can be associated to multiple workload templates. For example, two similar workload templates might offer the same service.

1 On the main navigation bar, click  *Catalog.*

2 Click *Legal Agreements*, then click *Create* to open a Create Legal Agreement dialog box.

3 In the *General* section of the dialog box, provide the following details for the legal agreement:

Name: Specify a name for the legal agreement. The same legal agreement can be associated to one or more workload templates, add-on applications or add-on services.

Version: Define a version of the legal agreement. This helps you track the changes in legal agreements from the software vendor, or from you as a service provider.

Agreement Text: Enter (or copy and paste) text that describes the terms, conditions, or legal agreements that must be accepted by a business service sponsor or the requestor of the add-on application or service. These terms must be accepted prior to adding a workload to a business service request.

You can use HTML markup in this field, perhaps copied from an HTML editor or existing HTML file. Plain text, HTML elements, and JavaScript are supported. CSS is not currently supported.

You can preview the result of your markup inside a Business Service Request when you edit a workload and view its Legal Agreement.

Creation Date: Displays the current date.

4 (Optional) You can associate the legal agreement with one or more workload templates:

4a Under *Associations*, click the *Workload Templates* tab.

4b Click *Add* to display the Add Workload Templates dialog box

4c Select the workload templates to add.

You can Shift-click and Ctrl-click to select multiple workload templates.

4d Click *OK* to add the selected workload templates to the *Workload Templates* list.

5 (Optional) You can associate the legal agreement with an add-on service or an add-on application.

5a Under *Associations*, click the *Add-Ons* tab.

5b Choose either *Applications* or *Services*, then click *Add* to display the list of available applications or services that you can choose to associate to this legal agreement.

5c On the Applications or Services dialog box, choose one or more applications or services to which you want to associate this legal agreement, click *OK*, then click *Save*.

12.7.2 Editing a Legal Agreement

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

To edit a legal agreement in the catalog:

1 On the main navigation bar, click  *Catalog.*

2 On the Legal Agreements tab, select an unaccepted legal agreement from the list, then click *Edit* to open a Configure Legal Agreement dialog box.

The following table lists the tasks and actions that you might want to perform as you configure a legal agreement in the Cloud Manager catalog.

Task	Steps
Change the legal agreement details	1. In the *General* section, modify the following legal agreement details: **Name:** Specify a name for the legal agreement. The same legal agreement can be associated to one or more workload templates, add-on applications or add-on services. **Version:** Define a version of the legal agreement. This helps you track the changes in legal agreements from the software vendor, or from you as a service provider. **Agreement Text:** Enter (or copy and paste) text that describes the terms, conditions, or legal agreements that must be accepted by a business service sponsor or the requestor of the add-on application or service. These terms must be accepted prior to adding a workload to a business service request. You can use HTML markup in this field, perhaps copied from an HTML editor or existing HTML file. Plain text, HTML elements, and JavaScript are supported. CSS is not currently supported. You can preview the result of your markup inside a Business Service Request when you edit a workload and view its Legal Agreement. **Creation Date:** Displays the date the legal agreement was created. You cannot change this setting.
Change how add-on applications or services are associated	You can change the association of the legal agreement with one or more workload templates: 1. Under *Associations*, click the *Workload Templates* tab. 2. Click *Add* to display the Add Workload Templates dialog box. 3. Select the workload templates to add. You can Shift-click and Ctrl-click to select multiple workload templates. 4. Click *OK*. You can change the association of the legal agreement with an add-on service or an add-on application. 1. Under *Associations*, click the *Add-Ons* tab. 2. Choose either *Applications* or *Services*, then click *Add* to display the list of available applications or services that you can choose to associate to this legal agreement. 3. On the *Applications* or *Services* dialog box, choose one or more applications or services to which you want to associate this legal agreement, click *OK*, then click *Save*.

12.8 Resource Pricing Packages

As a Cloud Administrator or Catalog Manager, you can create and edit a variety of resource packages (that is, the resources you configure and assign to a VM) with a different price points, applied discounts, and contract lengths. You can even assign a minimum contract length that must be fulfilled to receive special pricing.

- Section 12.8.1, “Creating a Resource Pricing Package,” on page 127
- Section 12.8.2, “Configuring a Resource Pricing Package,” on page 128

For more information about business contracts, see Chapter 15.11, “Understanding Business Service Contracts,” on page 184.

12.8.1 Creating a Resource Pricing Package

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

You can create resource pricing packages to apply special pricing and discounts for customers who agree to use customized workload resources for a specific contract period.When created, these packages are listed in the Cloud Manager catalog. From the catalog, you can edit a package that has not yet been accepted by a user as part of a workload, or you can view it to check its associations.

A single resource pricing package can be associated to multiple workload templates. For example, two similar workload templates might offer the same package. You can also associate a resource pricing package with all of the workload templates in a zone.

1. On the main navigation bar, click  *Catalog.*
2. Click *Resource Pricing Packagess*, then click *Create* to open a Create Resource Pricing Package dialog box.
3. In the *General* section of the dialog box, provide the following details for the resource pricing package:

   Description: Specify a description of the pricing package that can provide more information about the nature of the pricing package.

   Name: Specify a name for the resource pricing package. The same package can be associated to one or more workload templates.

   Minimum Contract: Choose a minimum contract length from the drop-down list. The user must agree to this contract length.

   Renewal Contract: Choose a duration period for a renewable contract for the resource pricing package. The user is notified about the expiration of the contract in time for renewal.
4. In the *Virtual Machine Resources* section:
   1. Customize the following settings to increase or decrease the workload resources.

      Number of CPUs: Select the minimum and maximum number of CPUs for a workload. The maximum allowed vCPUs cannot exceed the number supported by the hypervisor.

      You can also apply a discount to this resource, to give the customer an incentive to purchase the package.

      Memory: Select the minimum and maximum number megabytes (MB) of RAM to allocate to a workload.

You can also apply a discount to this resource, to give the customer an incentive to purchase the package.

Number of NICs: Select the number of network interface cards (NICs) to allocate to the workload.

You can also apply a discount to this resource, to give the customer an incentive to purchase the package.

Total Storage: Displays the amount of virtual storage (in GB) that you can designate as the minimum and maximum allowed for this resource pricing package.

You can also apply a discount to this resource, to give the customer an incentive to purchase the package.

2. (Optional) Select *Try these values* to open a worksheet dialog where you can experiment with different discount percentages and apply them to each resource. The list shows the effect of the discounts applied to each service level. You can click *Apply* to bring your tested values forward as the actual configured resource values.

5 (Optional) You can associate the resource pricing package with one or more workload templates:

5a Under *Associations*, click the *Workload Templates* tab.

5b Click *Add* to display the Add Workload Templates dialog box

5c Select the workload templates to add.

You can Shift-click and Ctrl-click to select multiple workload templates.

If you want to associate the resource pricing package with all workload templates in the catalog, select *Assign to all Workload Templates*.

5d Click *Save* to add the selected workload templates to the *Workload Templates* list.

NOTE: If any resource pricing packages are associated with a workload template, the business service owner is required to pick a package when configuring a workload. If you want to provide a way for the business service owner to opt out of picking a package, you need to create a package without restrictions or discounts. For example, you could create a package called “No Package” with no restrictions and no discounts.

12.8.2 Configuring a Resource Pricing Package

Roles that Can Perform This Task: Cloud Administrator, Catalog Manager

---

To configure a resource pricing package in the catalog:

1 On the main navigation bar, click  *Catalog.*

2 On the Legal Agreements tab, select a resource pricing package from the list, then click *Edit* to open an Edit Resource Pricing Package dialog box.

The following table lists the tasks and actions that you might want to perform as you configure a resource pricing package in the Cloud Manager catalog.

Task	Steps
Change the resource pricing package details	1. In the *General* section, modify the following resource pricing package details: **Name:** Specify a name for the resource pricing package. The same package can be associated to one or more workload templates. **Description:** Specify a description of the pricing package that can provide more information about the nature of the pricing package. **Minimum Contract:** Choose a minimum contract length from the drop-down list. The user must agree to this contract length. **Renewal Contract:** Choose a duration period for a renewable contract for the resource pricing package. The user is notified about the expiration of the contract in time for renewal. 2. In the *Virtual Machine Resources* section: a. Customize the following settings to increase or decrease the workload resources. **Number of CPUs:** Select the minimum and maximum number of CPUs for a workload. The maximum allowed vCPUs cannot exceed the number supported by the hypervisor. You can also apply a discount to this resource, to give the customer an incentive to purchase the package. **Memory:** Select the minimum and maximum number megabytes (MB) of RAM to allocate to a workload. You can also apply a discount to this resource, to give the customer an incentive to purchase the package. **Number of NICs:** Select the number of network interface cards (NICs) to allocate to the workload. You can also apply a discount to this resource, to give the customer an incentive to purchase the package. **Total Storage:** Displays the amount of virtual storage (in GB) that you can designate as the minimum and maximum allowed for this resource pricing package. You can also apply a discount to this resource, to give the customer an incentive to purchase the package. b. (Optional) Select *Try these values* to open a worksheet dialog where you can experiment with different discount percentages and apply them to each resource. The list shows the effect of the discounts applied to each service level. You can click *Apply* to bring your tested values forward as the actual configured resource values.

Task	Steps
Change how resource pricing packages are associated	You can associate the resource pricing package with one or more workload templates: 1. Under *Associations*, click the *Workload Templates* tab. 2. Click *Add* to display the Add Workload Templates dialog box. 3. Select the workload templates to add. You can Shift-click and Ctrl-click to select multiple workload templates. 4. Click *Save*. If you want to associate the resource pricing package with all workload templates in the catalog, select *Assign to all Workload Templates*. **NOTE:** If any resource pricing packages are associated with a workload template, the business service owner is required to pick a package when configuring a workload. If you want to provide a way for the business service owner to opt out of picking a package, you need to create a package without restrictions or discounts. For example, you could create a package called “No Package” with no restrictions and no discounts.

13 Managing Organizations

The following sections provide information to help you manage the organizations in your Cloud environment:

- Section 13.1, “Creating Organizations,” on page 131
- Section 13.2, “Customizing Organization Configuration Settings,” on page 133
- Section 13.3, “Creating Business Groups,” on page 134
- Section 13.4, “Assigning Resource Groups to Organizations and Business Groups,” on page 136
- Section 13.5, “Managing Organization Users, User Groups, and Roles,” on page 138

13.1 Creating Organizations

Roles that Can Perform This Task: Cloud Administrator

An organization represents a tenant to which you are offering Cloud services. Through the organization, you make resource group assignments that dictate the hosts, service levels, repositories, and networks available to the organization, and make workload template assignments that determine the types of business service workloads available to the organization.

After you create an organization, you (or an Organization Manager) can define the organization’s membership and assign roles such as Business Service Owner, Business Group Sponsor, and Organization Manager to those members. Membership and role assignments are covered in the next task, *Create Users and Groups*.

1. On the main navigation bar, click *Organizations,* click the *Organizations* tab, then click *Create*.
2. Provide the following details to define the organization:

   Name: Specify a unique name for the organization. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Description: Provide any additional information to identify the organization.

   Domains: If you want to enable users to self-register in this organization, specify the e-mail domains associated with the organization.

   Self-registration occurs when a valid LDAP user who does not have a Cloud Manager account first logs in. The user’s e-mail domain is compared to the e-mail domains defined for the organization. If it matches one of the e-mail domains, the user is added to organization’s Members list.

   You can associate one or more e-mail domains with the organization. To specify multiple e-mail domains, separate the names with commas (for example, `netiq.com,novell.com,attachmate.com`).

   For more information, see Section 11.2, “Enabling Automatic User Creation,” on page 105.

   Discount %: If you want a discount applied to all business services created by members of this organization, specify the discount percentage.

Auto Approval: When a user creates a business service request, the request goes through an approval workflow that includes both a Sponsor and an Administrator. The Sponsor is a member of the organization who provides the financial approval for the business service. The Administrator is a System user (such as yourself, another Cloud Administrator, or a Zone Administrator) who provides the resource capacity approval for the business service. You can use Auto Approval to bypass one or both of the approvals.

The organization inherits the Auto Approval settings from the Cloud Manager system settings (accessed through *Configuration* on the main navigation bar). To change the settings for the organization, click *Override*, then configure the settings as desired.

Logo: You can upload a logo file for the organization. Three formats are supported: PNG, JPG, and GIF. Any size is acceptable. Cloud Manager resizes the logo to a maximum of 216x216 pixels, maintaining the width-to-height proportions. For example, a 432x200 image would be resized to 216x100. The logo file is stored on the Cloud Manager Application Server.

To upload a file, mouse over *No Image*, then click *Upload New Image*. Browse for and select the image, then click *OK* to upload it to the Cloud Manager Application Server.

3 Add the workload templates that you want the organization to have access to.

You do not need to assign workload templates to the organization at this time. If you want to do it later, see Section 12.2.1, “Assigning Workload Templates to an Organization,” on page 116 when you are ready.

3a Under *Membership and Access*, click the *Workload Templates* tab.

3b Click *Add* to display the Add Workload Templates dialog box.

3c Select the workload templates.

You can Shift-click and Ctrl-click to select multiple workload templates.

3d Click *OK* to add the selected workload templates to the *Workload Templates* list.

4 Add the resource groups that you want the organization to have access to.

You do not need to assign resource groups to the organization at this time. If you want to do it later, see “Assigning Resource Groups to Organizations and Business Groups” on page 136 when you are ready.

4a Under *Membership and Access*, click the *Resource Groups* tab.

4b Click *Add* to display the Add Resource Groups dialog box.

4c Select the resource groups you want to add.

You can Shift-click and Ctrl-click to select multiple groups.

4d Click *OK* to add the selected resource groups to the *Resource Groups* list.

5 Add the networks that you want the organization to have access to.

The available networks are determined by the VM hosts included in the resources groups. However, to enable you to provide isolated networks for two or more organizations that share the same resource group, the networks from a resource group are not automatically assigned to an organization when you add the resource group. Instead, you must separately add the networks you want assigned to the organization.

5a Under *Membership and Access*, click the *Networks* tab.

5b Click *Add* to display the Add Networks dialog box.

5c Select the networks.

You can Shift-click and Ctrl-click to select multiple networks.

5d Click *OK* to add the selected networks to the *Networks* list.

6 Add organization members.

6a Click *Save* to create the organization.

You can only add users after the organization has been saved.

6b Refer to "Manually Creating System and Organization Users" on page 101 for details about creating users and adding them to an organization. Or refer to "Manually Importing Organization Users from LDAP" on page 103 for details about importing users from an LDAP source into the organization.

7 Assign roles for the organization.

Users must be given roles in order to do anything in the organization. There are six roles that apply to an organization: Approver, Build Administrator, Business Group Viewer, Business Service Owner, Organization Manager, and Sponsor.

Role assignments at the organization level are inherited by the organization's business groups. For example, if you give a user the Business Service Owner role for an organization, the user can create business services for any business group in the organization. If you want to limit the user to a role in specific business group, you must make the role assignment in the business group.

7a Click the *Users* tab, then click the role (*Approver*, *Build Administrator*, *Business Group Viewer*, *Business Service Owner, Organization Manager*, or *Sponsor*) that you want to assign to a user.

7b Click *Add*.

Depending on the role that you are adding, the selection dialog box can contain two lists: *Members* and *System Users*. The *Members* list includes all members of the organization and the *System Users* list includes all Cloud Manager System users.

7c Select the users you want to add, then click *OK*.

You can Shift-click and Ctrl-click to select multiple users.

8 Click *Save* to add the organization to the *Organizations* list.

13.2 Customizing Organization Configuration Settings

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

The Organization Configuration settings determine if organization members can see all workload costs and settings when creating and managing business services.

To customize the settings:

1 On the main navigation bar, click *Organizations.*

2 On the *Organizations* tab, select the target organization, then click *Edit*.

3 In the upper-right corner of the dialog box, click the icon to display the Organization Configuration dialog box.

4 Configure the following settings:

Organization Costs: Select whether the organization's business service costs are hidden from or shown to the organization's members. This applies to all Organization members regardless of their assigned roles.

You can override this setting on a business group. For example, if you show business service costs for the organization, you can hide costs for a specific business groups. Organization users can see the costs associated with all of the organization's business services with the exception of business services created by the one business group.

In some cases, you might want to hide costs from some users but not from others. To achieve this, you can hide the costs at the organization or business group level, but then enable the *Always show costs* option on the individual user accounts.

Workload Dialog: A workload contains *Windows Settings*, *Windows Licensing*, *Linux Settings*, and *Networks* tabs that must be configured when requesting a business service. These tabs can be hidden, in which case a Build Administrator or Cloud Administrator must provide the information when completing the pre-build configuration tasks for the business service.

The organization inherits the Workload Dialog settings configured at the system level. If you want to override the system settings to apply different settings for the organization, select *Override System Setting*, then enable or disable the tabs as desired.

To revert the settings to the system settings, deselect *Override System Setting.*

5 Click *OK*.

13.3 Creating Business Groups

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

A business group represents a unit within an organization, such as a department or cost center, with which business services are associated. An organization can have one or more business groups.

A business group can be assigned all of an organization's resources or only some of the resources. When a business service is created for a business group, it uses only the assigned resources. Multiple business groups can be assigned the same resources, which means that the resources become shared resources.

1 On the main navigation bar, click *Organizations*.

2 Click the *Business Groups* tab, then click *Create*.

3 Provide the following details to define the business group:

Name: Specify a name for the group. The name should be different than any other business group name.

Organization: Select the organization for the business group. The organization assignment cannot be changed after the business group is created.

Description: Provide any additional information to identify the business group.

Auto Approval: Business service requests require both a Sponsor approval and an Administrator approval. The Sponsor approval is a financial check, while the Administrator approval is a resource capacity check. You can use Auto Approval to bypass one or both of the approvals.

Sponsor is selected by default. If you don't want automatic Sponsor approval, you must add sponsors to the group (see Step 4).

Select *Administrator* to automatically grant Administrator approval for the group's business services.

Costs: The business group inherits the *Costs* setting from its organization. To change the setting for the business group, click *Override*, then configure the setting as desired. *Show* allows group members to see cost information for workloads. *Hide* to prevent group members from seeing cost information.

4 Assign roles for the business group.

There are three roles that apply to a business group: Business Group Viewer, Business Service Owner, and Sponsor. By default, users assigned these roles at the organization also have these same roles in the business group.

4a Click the *Users* tab, then click the role (*Business Group Viewer*, *Business Service Owner,* or *Sponsor*) that you want to assign to a user.

4b Click *Add*.

Depending on the role that you are adding, the selection dialog box can contain two lists: *Members* and *System Users*. The *Members* list includes all members of the organization and the *System Users* list includes all Cloud Manager System users.

4c Select the users you want to add, then click *OK*.

You can Shift-click and Ctrl-click to select multiple users.

5 Add the workload templates that you want the business group to have access to.

You do not need to assign workload templates to the business group at this time. If you want to do it later, see Section 12.2.2, “Assigning Workload Templates to a Business Group,” on page 116 when you are ready.

5a Under *Membership and Access*, click the *Workload Templates* tab.

5b Click *Add* to display the Add Workload Templates dialog box.

The list displays the organization’s workload templates. A business group is limited to the workload templates assigned to its organization.

5c Select the workload templates.

You can Shift-click and Ctrl-click to select multiple workload templates.

5d Click *OK* to add the selected workload templates to the *Workload Templates* list.

6 Add the resource groups you want the business group to have access to.

You do not need to assign resource groups to the organization at this time. If you want to do it later, see “Assigning Resource Groups to Organizations and Business Groups” on page 136 when you are ready.

6a Under *Membership and Access*, click the *Resource Groups* tab.

6b Click *Add* to display the Add Resource Groups dialog box.

The list displays the organization’s resource groups. A business group is limited to the resource groups assigned to its organization.

6c Select the resource groups you want to add.

You can Shift-click and Ctrl-click to select multiple resource groups.

6d Click *OK* to add the selected resource groups to the *Resource Groups* list.

7 Add the networks that you want the business group to have access to.

The available networks are determined by the VM hosts included in the resource groups. However, to enable you to provide isolated networks for business groups that share the same resource group, the networks from a resource group are not automatically assigned to a business group when you add the resource group. Instead, you must separately add the networks you want assigned to the business group.

7a Under *Membership and Access*, click the *Networks* tab.

7b Click *Add* to display the Add Networks dialog box.

7c Select the networks.

You can Shift-click and Ctrl-click to select multiple networks.

7d Click *OK* to add the selected networks to the *Networks* list.

8 Click *Save*.

13.4 Assigning Resource Groups to Organizations and Business Groups

Resource groups must be assigned to an organization in order for the organization's business services to use those resources. After you assign resource groups to an organization, you or the Organization Manager can make them available to business groups within the organization.

The following steps assume that the resource groups you want to assign already exist. If they do not, see "Creating, Modifying, and Deleting Resource Groups" on page 141.

- Section 13.4.1, "Assigning Resource Groups to an Organization," on page 136
- Section 13.4.2, "Assigning Resource Groups to a Business Group," on page 137
- Section 13.4.3, "Removing Resource Group Assignments from an Organization," on page 138
- Section 13.4.4, "Removing Resource Group Assignments from a Business Group," on page 138

13.4.1 Assigning Resource Groups to an Organization

Roles that Can Perform This Task: Cloud Administrator

1 On the main navigation bar, click *Organizations.*

2 On the *Organizations* tab, select the target organization, then click *Edit*.

3 Add the resource groups you want the organization to have access to:

3a Under *Membership and Access*, click the *Resource Groups* tab.

3b Click *Add* to display the Add Resource Groups dialog box.

3c Select the resource groups to add.

You can Shift-click and Ctrl-click to select multiple groups.

3d Click *OK* to add the selected resource groups to the *Resource Groups* list.

4 Add the networks that you want the organization to have access to.

The resource group's networks are determined by its VM hosts. To enable you to provide isolated networks for organizations that share the same resource group, the networks are not automatically added when you add the resource group. You must explicitly add the networks you want assigned to the organization.

4a Under *Membership and Access*, click the *Networks* tab.

4b Click *Add* to display the Add Networks dialog box.

4c Select the networks to add.

You can Shift-click and Ctrl-click to select multiple networks.

4d Click *OK* to add the selected networks to the *Networks* list.

5 When you have finished adding resource groups, click *Save* to save the changes.

13.4.2 Assigning Resource Groups to a Business Group

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

A business group does not automatically inherit the resource groups assigned to its organization. Any of the organization's resource groups that you want to be available to the business group must be assigned to it.

1 On the main navigation bar, click *Organizations*, then click the *Business Groups* tab.

2 In the *Business Groups* list, select the target business group, then click *Edit*.

3 Add the resource groups you want the business group to have access to:

3a Under *Membership and Access*, click the *Resource Groups* tab.

3b Click *Add* to display the Add Resource Groups dialog box.

3c Select the resource groups to add.

You can Shift-click and Ctrl-click to select multiple groups.

3d Click *OK* to add the selected resource groups to the *Resource Groups* list.

4 Add the networks that you want the business group to have access to.

The resource group's networks are determined by its VM hosts. To enable you to provide isolated networks for business groups that share the same resource group, the networks are not automatically added when you add the resource group. You must explicitly add the networks you want assigned to the business group.

4a Under *Membership and Access*, click the *Networks* tab.

4b Click *Add* to display the Add Networks dialog box.

4c Select the networks to add.

You can Shift-click and Ctrl-click to select multiple networks.

4d Click *OK* to add the selected networks to the *Networks* list.

5 When you have finished adding resource groups, click *Save* to save the changes to the business group.

13.4.3 Removing Resource Group Assignments from an Organization

Roles that Can Perform This Task: Cloud Administrator

You can remove a resource group from an organization only if the resource group is not hosting any deployed business services from the organization's business groups.

Removing a resource group assignment from an organization also removes any assignments from its business groups.

1. On the main navigation bar, click *Organizations*.
2. In the *Organizations* list, select the organization from which you want to remove the resource group assignment, then click *Edit*.
3. Click the *Resource Groups* tab.
4. Select the resource group to be removed, click *Remove*, then click *Yes* to confirm the removal.

   You can Shift-click or Ctrl-click to select multiple resource groups.
5. When you have finished removing resource groups, click *Save* to save the changes to the organization.

13.4.4 Removing Resource Group Assignments from a Business Group

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

You can remove a resource group from a business group only if the resource group is not hosting any of the business group's deployed business services.

1. On the main navigation bar, click *Organizations*.
2. Click the *Business Groups* tab, select the business group from which you want to remove the resource group assignment, then click *Edit*.
3. Click the *Resource Groups* tab.
4. Select the resource group to be removed, click *Remove*, then click *Yes* to confirm the removal.

   You can Shift-click or Ctrl-click to select multiple resource groups.
5. When you have finished removing resource groups, click *Save* to save the changes to the business group.

13.5 Managing Organization Users, User Groups, and Roles

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

You can add users and user groups to an organization and then assign roles to the users and groups so that they can perform specific functions within the organization.

Users, user groups, and roles are covered in Chapter 11, “Setting Up and Managing Users,” on page 101.

14 Managing Resources

The following sections provide information to help you manage the resources in your Cloud environment:

- Section 14.1, “Creating, Modifying, and Deleting Resource Groups,” on page 141
- Section 14.2, “Creating, Modifying, and Deleting Service Levels,” on page 144
- Section 14.3, “Assigning Service Levels to Resource Groups,” on page 148
- Section 14.4, “Editing Network Configurations,” on page 149
- Section 14.5, “Monitoring Resource Capacity,” on page 154

14.1 Creating, Modifying, and Deleting Resource Groups

A resource group defines a set of VM hosts that an organization can use for its business services. In addition to the VM hosts, the resource group includes one or more service levels that define the cost of the host resources (vCPUs, memory, storage, and networks) and the service objectives (availability, support response time, and so forth).

- Section 14.1.1, “Creating Resource Groups,” on page 141
- Section 14.1.2, “Modifying Resource Groups,” on page 143
- Section 14.1.3, “Deleting Resource Groups,” on page 143

14.1.1 Creating Resource Groups

Roles that Can Perform This Task: Cloud Administrator, Zone Administrator (create resource groups and add hosts only; cannot add service levels or assign resource groups to organizations)

1. On the main navigation bar, click  *Resources.*
2. Click the *Resource Groups* tab, then click *Create*.
3. If you are the Zone Administrator for multiple zones, the Select Zone dialog box is displayed. Select the zone that contains the resources you are grouping, then click *OK* to display the Create Resource Group dialog box.
4. In the *General* fields, provide the following information for the resource group:

   Name: Specify a unique name for the group. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Zone: Displays the zone whose hosts you can add to the group. You cannot change this setting.

   Hypervisor: Select the hypervisor technology for the group’s hosts. You can add only those hosts that meet the hypervisor criteria.

Workload Repository: The *Default* setting causes a provisioned workload to be stored in the same repository as the VM template used to create it. If you want workloads provisioned to this resource group to be stored in a different shared repository, you must add hosts to the group (see Step 5), then come back and select the shared repository for the workloads. The *Workload Repository* list is populated only after you add hosts to the resource group.

Group Type: This applies only if VMware vSphere is the selected hypervisor. Select *Host* if you want the resource group to use hosts and host clusters. Select *Resource Pool* if you want the resource group to use a resource pool.

Resource Pool: If you specified *Resource Pool* as the group type, select the resource pool to include in the group.

Description: Provide any additional information for the resource group.

5 If the group type is *Host*, add hosts to the group:

5a Under *Associations*, click the *Hosts* tab.

5b Click *Add* to display the Add Hosts dialog box.

The list displays all available hosts and host clusters in the zone that meet the selected hypervisor criteria. Hosts that are already assigned to another resource group are not displayed.

5c Select the hosts.

You can Shift-click and Ctrl-click to select multiple hosts.

5d Click *OK* to add the selected hosts to the *Hosts* list.

6 Add service levels to the group:

6a Under *Associations*, click the *Service Level* tab.

6b Click *Add* to display the Add Service Levels dialog box.

6c Select the service levels.

You can Shift-click and Ctrl-click to select multiple service levels.

6d Click *OK* to add the selected service levels to the *Service Levels* list.

7 Ignore the *Networks* tab.

The *Networks* tab shows the networks associated with the hosts you added to the group. The list is view-only so you can't make any changes. However, the list is not generated until you save the resource group. If you want to see the networks at this time, click *Save*, double-click the resource group to open it again, then click the *Networks* tab.

8 Specify the organizations that can use the resource group:

8a Under *Associations*, click the *Organizations* tab.

8b Click *Add* to display the Add Organizations dialog box.

8c Select the organizations to which you want to assign the resource group.

You can Shift-click and Ctrl-click to select multiple hosts.

8d Click *OK* to add the selected organizations to the *Organizations* list.

IMPORTANT: The resource group is added to the organization, but its networks are not made available to the organization. To make the networks available, edit the organization and add the resource group's networks to the *Networks* list.

---

9 Click *Save*.

14.1.2 Modifying Resource Groups

Roles that Can Perform This Task: Cloud Administrator, Zone Administrator (add and remove hosts only)

You can modify a resource group to add or remove hosts, to add or remove service levels, and to change the organization assignments.

Be aware of the following when you remove hosts and service levels from a resource group and a resource group from an organization:

- You can remove a host at any time. Any business service workloads running on the host continue to run until they are cycled (stopped, then started) in the Cloud Manager console. At that point, they are moved to another host in the resource group.
- You cannot remove a service level if it is associated with a business service that is deployed to the resource group.
- You cannot remove a resource group assignment from an organization if the organization has business services deployed to the resource group. In addition, removing a resource group assignment from an organization also removes any assignments from its business groups.

To modify a resource group:

1. On the main navigation bar, click *Resources.*
2. Click the *Resource Groups* tab, select the resource group you want to modify, then click *Edit*.
3. Make the desired changes to the resource group, then click *OK* to save the changes.

   For a description of the resource group settings, see “Creating, Modifying, and Deleting Resource Groups” on page 141 or click the icon in the Edit Resource Group dialog box.

14.1.3 Deleting Resource Groups

Roles that Can Perform This Task: Cloud Administrator, Zone Administrator

You can delete a resource group only if the resource group is not hosting any deployed business services.

1. On the main navigation bar, click *Resources.*
2. On the *Resource Groups* tab, select the resource group to delete, then click *Delete*.

   If the *Delete* action is not available, the resource group is hosting deployed business services and cannot be deleted.
3. Click *Yes* to confirm the deletion.

14.2 Creating, Modifying, and Deleting Service Levels

A service level defines the monthly cost for each type of host resource (vCPUs, memory, storage, and networks). For example, you might set the cost of one vCPU at $25 per month. If a workload requires two vCPUs, $50 is added to the monthly cost of the workload.

A service level can also include service objectives. Objectives typically define measurable behaviors such as host availability (uptime) or support response time and have a cost associated with them. Any service objective costs are added to the monthly cost of a workload that is deployed in the resource group.

The following sections provide instructions for managing service levels and service level objectives:

- Section 14.2.1, “Creating Service Levels,” on page 144
- Section 14.2.2, “Modifying Service Levels,” on page 145
- Section 14.2.3, “Deleting Service Levels,” on page 146
- Section 14.2.4, “Creating Service Level Objectives,” on page 146
- Section 14.2.5, “Modifying Service Level Objectives,” on page 147
- Section 14.2.6, “Deleting Service Level Objectives,” on page 148

14.2.1 Creating Service Levels

Roles that Can Perform This Task: Cloud Administrator

1. On the main navigation bar, click *Resources.*
2. Click the *Service Levels* tab, then click *Create*.
3. In the *General* section, provide the following details for the service level:

   Name: Specify a unique name for the service level. This name is displayed in business service workloads.

   You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

   Creation Date: Displays the current date.

   Description: Provide any additional information for the service level.
4. In the *Monthly Resource Costs* fields, define the cost (per month) to use the host resources:

   vCPU: Specify the cost per virtual CPU.

   Memory: Specify the cost per megabyte (MB) of memory.

   Disk: Specify the cost per gigabyte (GB) of disk space.

   Network: Specify the cost per network interface card.
5. (Optional) Add objectives to the service level.
   - **5a** Under *Associations*, click the *Service Level Objectives* tab.
   - **5b** Click *Add* to display the Add Service Level Objectives dialog box.
   - **5c** Select the objectives to add.

     You can Shift-click and Ctrl-click to select multiple objectives.

If you have not yet created the objectives you want, see “Creating Service Level Objectives” on page 146

5d Click *OK* to add the selected objectives to the *Service Level Objectives* list.

6 Assign the service level to the appropriate resource groups:

6a Under *Associations*, click the *Resource Groups* tab.

6b Click *Add* to display the Add Resource Groups dialog box.

6c Select the groups to add.

You can Shift-click and Ctrl-click to select multiple groups.

6d Click *OK*.

7 Click *Save*.

14.2.2 Modifying Service Levels

Roles that Can Perform This Task: Cloud Administrator

You can modify a service level to change the resource costs and objectives. Changing a service level can impact the cost of business services currently using the service level.

1 On the main navigation bar, click  *Resources*, then click the *Service Levels* tab.

2 Select the service level you want to modify, then click *Edit.*

3 In the *General* section, modify the details for the service level:

Name: Specify a unique name for the service level. This name is displayed in business service workloads.

You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The maximum length is 110 characters.

Creation Date: Displays the current date.

Description: Provide any additional information for the service level.

4 In the *Monthly Resource Costs* fields, define the cost (per month) to use the host resources:

vCPU: Specify the cost per virtual CPU.

Memory: Specify the cost per megabyte (MB) of memory.

Disk: Specify the cost per gigabyte (GB) of disk space.

Network: Specify the cost per network interface card.

5 Add objectives to the service level:

5a Under *Associations*, click the *Service Level Objectives* tab.

5b Click *Add* to display the Add Service Level Objectives dialog box.

5c Select the objectives to add.

You can Shift-click and Ctrl-click to select multiple objectives.

If you have not yet created the objectives you want, see “Creating Service Level Objectives” on page 146

5d Click *OK* to add the selected objectives to the *Service Level Objectives* list.

6 Remove objectives from the service level:

6a Under *Associations*, click the *Service Level Objectives* tab.

6b In the *Service Level Objectives* list, select the objectives you want to remove.

You can Shift-click and Ctrl-click to select multiple objectives.

6c Click *Remove*.

7 Assign the service level to resource groups:

7a Under *Associations*, click the *Resource Groups* tab.

7b Click *Add* to display the Add Resource Groups dialog box.

7c Select the groups to add.

You can Shift-click and Ctrl-click to select multiple groups.

7d Click *OK*.

8 Remove the service level from resource groups:

8a Under *Associations*, click the *Resource Groups* tab.

8b Select the resource groups from which to remove the service level.

You can Shift-click and Ctrl-click to select multiple groups.

8c Click *Remove*.

9 Click *Save* to save your changes to the service level.

14.2.3 Deleting Service Levels

Roles that Can Perform This Task: Cloud Administrator

If a service level is associated with a business service, you cannot delete the service level.

1 On the main navigation bar, click *Resources*, then click the *Service Levels* tab.

2 Select the service level you want to delete, then click *Delete*.

3 Click *Yes* to confirm the deletion.

If the service level is associated with a workload, you receive a message informing you that the service level cannot be deleted.

14.2.4 Creating Service Level Objectives

Roles that Can Perform This Task: Cloud Administrator

Service level objectives define measurable characteristics such as availability, throughput, frequency, response time, and quality. Each of these characteristics typically has multiple objectives that identify a different level of service.

For example, you might define three availability objectives (97%, 98%, and 99%) and associate them with different service levels (Silver, Gold, and Platinum). By associating different costs with the objectives, you can establish the desired cost structure for your service levels.

1. On the main navigation bar, click *Resources.*
2. Click the *Service Levels* tab, click *Service Level Objectives*, then click *Create* to display the Create Service Level Objective dialog box.
3. Provide the following details:

   Name: Specify a name for the service level objective. The name should be different than any other objective name.This name not only appears to administrators but also to business service requestors when they configure workloads with service levels that include the objective.

   Monthly Cost: Specify the cost associated with the objective. If the objective does not have a cost, leave the field empty.

   Description: Provide optional text to further identify the service level objective.

   Creation Date: Displays the current date.

   Objective Type: If this objective represents workload availability, select *Availability*. Otherwise, select *General*.

   Value: If the objective type is *Availability*, specify the target availability as a percentage (for example, 99.9). If the objective type is *General*, specify an appropriate objective value.
4. Click *Save*.

14.2.5 Modifying Service Level Objectives

Roles that Can Perform This Task: Cloud Administrator

1. On the main navigation bar, click *Resources.*
2. Click the *Service Levels* tab, then click *Service Level Objectives*.
3. Select the objective you want to modify, then click *Edit*.
4. Modify the following details:

   Name: Specify a name for the service level objective. The name should be different than any other objective name.This name not only appears to administrators but also to business service requestors when they configure workloads with service levels that include the objective.

   Monthly Cost: Specify the cost associated with the objective. If the objective does not have a cost, leave the field empty.

   Description: Provide optional text to further identify the service level objective.

   Creation Date: Displays the current date.

   Objective Type: If this objective represents workload availability, select *Availability*. Otherwise, select *General*.

   Value: If the objective type is *Availability*, specify the target availability as a percentage (for example, 99.9). If the objective type is *General*, specify an appropriate objective value.
5. Click *Save*.

14.2.6 Deleting Service Level Objectives

Roles that Can Perform This Task: Cloud Administrator

You cannot delete a service level objective that is associated with a service level. Remove the objective from any service levels before deleting it.

1. On the main navigation bar, click *Resources.*
2. Click the *Service Levels* tab, then click *Service Level Objectives*.
3. Select the service level you want to delete, then click *Delete*.
4. Click *Yes* to confirm the deletion.

   If the objective is associated with a service level, you receive a message informing you of the association.

14.3 Assigning Service Levels to Resource Groups

A resource group identifies a collection of VM hosts to which workloads can be deployed. However, a resource group does not includes any costs associated with running workloads on the hosts. A resource group also does not include any service objectives for the workloads (such as host availability or support response time). The resource costs and service objectives are applied to resource groups through the assignment of service levels.

- Section 14.3.1, “Assigning Service Levels,” on page 148
- Section 14.3.2, “Removing Service Levels,” on page 148

14.3.1 Assigning Service Levels

Roles that Can Perform This Task: Cloud Administrator

1. On the main navigation bar, click *Resources.*
2. Click the *Resource Groups* tab, select the resource group to which you want to assign a service level, then click *Edit*.
3. Under *Associations*, click the *Service Level* tab.
4. Click *Add* to display the Add Service Levels dialog box.
5. Select the service level.

   You can Shift-click and Ctrl-click to select multiple service levels.
6. Click *OK* to add the selected service level to the *Service Levels* list.
7. Click *Save* to save the changes to the resource group.

14.3.2 Removing Service Levels

Roles that Can Perform This Task: Cloud Administrator

You cannot remove service levels that are associated with business services deployed in the resource group.

1. On the main navigation bar, click *Resources.*
2. Click the *Resource Groups* tab, select the resource group from which you want to remove the service level, then click *Edit*.
3. Under *Associations*, click the *Service Level* tab.
4. Select the service level to remove.

   You can Shift-click and Ctrl-click to select multiple service levels.
5. Click *Remove* to remove the selected service level from the *Service Levels* list.
6. Click *Save* to save the changes to the resource group.

   If the service level is associated with a business service that is deployed in the resource group, you receive a message stating that the service level cannot be removed.

14.4 Editing Network Configurations

The *Networks* list shows the networks (that is, the virtual LANs) in your Cloud Manager zones. For each network in the list, the following information is provided:

- **Name:** The network name. This is not configurable.
- **Zone:** The Cloud Manager zone where the network resides.
- **Configured By:** The role assigned to configure the network. This can be a *Business Service Requestor*, who can assign a static or DHCP-designated IP address, or a *Cloud Provider*, who can manually assign a static IP address or a DHCP-designated address, or who can delegate automatic IP address assignment through IPAM (IP address management).
- **Description:** The description of the network, as defined by the Cloud Administrator.

In addition to summarizing the network information, the list lets you select a network to view and edit its details. This section includes information about editing networks already discovered in your Cloud Manager zones.

- Section 14.4.1, “Editing the General Network Configuration,” on page 149
- Section 14.4.2, “Editing Network Associations,” on page 150
- Section 14.4.3, “Editing a Network’s IP Pool,” on page 151

14.4.1 Editing the General Network Configuration

Roles that Can Perform This Task: Cloud Administrator

1. On the main navigation bar, click *Resources.*
2. Click the *Networks* tab.
3. Select the network you want to modify, then click *Edit*.
4. On the Edit Network dialog, select the General tab.
5. In the *Network* area of the dialog, only the *Description* field is editable. Provide any additional information about the network, then click *Save* if you have no other changes to make to the network configuration.

6 In the *Network Assignment Method* fields, provide the following information for the network:

Network Configured By: You can specify how a NIC placed on the network is to be configured. It can be configured exclusively by you (the Cloud Provider), by a Business Service Requestor, or by using the network configuration defaults (defined in the *System Configuration* menu).

- If you specify that the network is to be configured according to system defaults (*Use System Defaults*), the network uses the settings defined in the *Configuration*  > *Networks* menu).
- If you specify that the network can be configured by a business service requestor, Cloud Manager locks the field to allow a configuration setting where the requestor can assign a static IP or DHCP-generated IP address to a workload in a business service.
- If you specify that the network must be configured exclusively by the cloud provider, this field provides some configuration options:
  - **Static (Manual):** By default, the NIC must be configured with a static IP address provided by you.
  - **DHCP (Manual):** This setting allows the automatic assignment of a DHCP address from the network to the network interface.
  - **IPAM system (Automatic):** You allow the NICs on the network to be automatically configured with an IP address assigned by an IP address management system (IPAM).

NOTE: For information about integrating the Nixu IPAM service with Cloud Manager, see the *NetIQ Cloud Manager 2.2.2 Procedures Guide* (https://www.netiq.com/documentation/cloudmanager222/ncm222_procedures/data/b12c23uw.html).

For information about setting up Cloud Manager native IPAM, see Section 14.4.3, “Editing a Network’s IP Pool,” on page 151.

The option to request a public address allows the user another, secondary address in addition to the primary NIC address. This secondary address is not provided by Cloud Manager. If the check box is selected when the pre-build stage of the workflow occurs, an administrator must either enter the configured public IP address in the adjacent field or override the request by deselecting the check box.

7 If you chose either the *Static* or the *IPAM* options as a Cloud Provider, in the *Network Default* fields, you can provide the *NetMask*, *Gateway*, *DNS Servers*, and *Search Domains* values for this network.

14.4.2 Editing Network Associations

Roles that Can Perform This Task: Cloud Administrator

---

1 On the main navigation bar, click  *Resources.*

2 Click the *Networks* tab.

3 Select the network you want to modify, then click *Edit*.

4 On the Edit Network dialog box, click the Associations tab.

5 (Optional) Enable organizations to use the network.

5a Click the *Organizations* tab.

5b Click *Add* to display the Add Organizations dialog box.

5c Select the organizations to which you want to assign the NIC.

You can Shift-click and Ctrl-click to select multiple NICs.

5d Click *OK* to add the selected organizations to the *Organizations* list.

NOTE: You can select *Associate with all Organizations and Business Groups* if you want to associate this NIC with all business groups and all organizations that exist in the Cloud Manager environment.

Only those NICs associated directly to a workload will be available for that workload.

---

6 (Optional) Remove an organization's access to the network.

6a Click the *Organizations* tab.

6b In the *Organizations* list, select the organization to remove.

You can Shift-click and Ctrl-click to select multiple organizations.

6c Click *Remove*.

7 (Optional) Enable business groups to use the network.

7a Click the *Business Groups* tab.

7b Click *Add* to display the Add Business Groups dialog box.

7c Select the business groups to which you want to assign the NIC.

You can Shift-click and Ctrl-click to select multiple NICs.

7d Click *OK* to add the selected business groups to the *Business Groups* list.

8 Remove a business group's access to the network.

8a Click the *Business Groups* tab.

8b In the *Business Groups* list, select the business group to remove.

You can Shift-click and Ctrl-click to select multiple business groups.

8c Click *Remove*.

14.4.3 Editing a Network's IP Pool

If the network is to be configured by you, the cloud provider, and you plan to use Cloud Manager's native IP address management system, you can specify values for the *NetMask* and *Gateway* network defaults on the *General* tab. This pre-populates the network with a pool of IP addresses that the system can assign to specific NICs on the network.

NOTE: You still need to specify the domain settings of the network, even if network defaults pre-populate this IP pool. See Step 5b below for more information.

If, however, the network defaults are not specified, and the *Network Assignment Method* is set to be configured by the *Cloud Provider* and the source as *IPAM*, you can still manually define a range for a pool of IP addresses, along with an assigned hostname of your choosing, and, if Linux, the domain name.

When a range of IP addresses is defined and listed as a pool, you can reserve one or more of those addresses so that they are not assigned to a NIC by the IPAM system.

The following steps describe how to edit a network IP pool:

1 On the main navigation bar, click *Resources.*

2 Click the *Networks* tab.

3 Select the network you want to modify, then click *Edit*.

4 On the Edit Network dialog box, click the Network tab.

5 (Optional) Specify an IP Range for use by the network.

5a Specify the IP Range:

Start: Specify the IP address you want to use for the first (starting) address of the pool range. The first useable address usually begins with *xxx.xxx.xxx.*1

End: Specify the IP address you want to use for the last (ending) address of the pool range. The last useable address is determined by the gateway prefix, if available.

5b Specify the Domain Settings:

Hostname Pattern: Specify a hostname that can help to identify the VM assigned to the IP. Windows VMs can have a hostname with a maximum of 15 characters. Linux VMs can have a hostname with a maximum of 63 characters.

By default, the field is populated with a optional, tokenized string value. Using a token is optional. The tokens substitute known values that can help you identify VMs in the pool according to a corresponding Cloud Manager object. For information about these tokens, see Hostname Assignment in Cloud Manager IPAM (page 152).

Domain Name: Specify the default domain name to be used by Linux VMs on this network. This should be an easily recognizable and memorizable name.

6 (Optional) Reserve an IP address or range of addresses for use by the network.

6a Select the available IP address or addresses you want reserved in the IP pool.

You can Shift-click and Ctrl-click to select multiple addresses.

6b Click *Reserve* to reserve the formerly available addresses to the IP pool.

NOTE: It is possible to include the IP Addresses of newly-discovered VMs in a Cloud Manager Orchestration zone in the IP address pool. For more information, see

---

Hostname Assignment in Cloud Manager IPAM

For IP pools, Cloud Manager requires you to create host names within the following parameters:

- Only valid characters can be included in the name: numbers, letters, and the dash (that is, the hyphen [-]) character.
- Windows VMs can have a hostname with a maximum of 15 characters.
- Linux VMs can have a hostname with a maximum of 63 characters.

In the *Host Name* field, you can specify tokenized strings that Cloud Manager later substitutes with information from its system:

Token	Notes
`%Short_IP%`	In an IP address of the format `a.b.c.d`, the short IP is the last part of the address, which is `d`
`%Long_IP%`	In an IP address of the format `a.b.c.d`, the long IP is the last two parts of the address, separated by a dash: `c-d`
`%Org_Name%`	◆ “All lowercase” token is `%org_name%` ◆ “All uppercase” token is `%ORG_NAME%`

Token	Notes
%BG_Name%	◆ “All lowercase” token is %bg_name% ◆ “All uppercase” token is %BG_NAME%
%BS_Name%	◆ “All lowercase” token is %bs_name% ◆ “All uppercase” token is %BS_NAME%
%WL_Name%	◆ “All lowercase” token is %wl_name% ◆ “All uppercase” token is %WL_NAME%
%Org_ID%	The value of the %Org_ID% token might collide with the IP value (see %Short_IP% or %Long_IP% above). Because Cloud Manager truncates the hostname after token replacement, keeping the wrong instance of the number is possible.
%BG_ID%	The value of the %BG_ID% token might collide with the IP value (see %Short_IP% or %Long_IP% above). Because Cloud Manager truncates the hostname after token replacement, keeping the wrong instance of the number is possible.
%BS_ID%	The value of the %BS_ID% token might collide with the IP value (see %Short_IP% or %Long_IP% above). Because Cloud Manager truncates the hostname after token replacement, keeping the wrong instance of the number is possible.
%WL_ID%	The value of the %WL_ID% token might collide with the IP value (see %Short_IP% or %Long_IP% above). Because Cloud Manager truncates the hostname after token replacement, keeping the wrong instance of the number is possible.

Cloud Manager analyzes either a token (or a combination of tokens) or a string you have specified and then truncates that string or substituted token value to fit the character length limitation of the hostname.

In this truncation process, Cloud Manager follows hierarchical rules as it brings the name into character-length compliance. That is, if the conditions of the first rule are satisfied, it executes the next rule, and so on, until compliance is reached.

IMPORTANT: Cloud Manager considers the IP address forms a critical part of hostname, so if you specify one of the IP address forms in a token, Cloud Manager identifies it as a “string to keep” in the truncated name:, so the %Short_IP% value or the %Long_IP% value is a “string to keep.” Any other strings or tokens surrounding the “string to keep” are considered either a prefix or a suffix to this string:

Prefix*StringToKeep*Suffix

If you specify both the Short_IP and Long_IP tokens in the field, the Long_IP token is retained and the Short_IP token is available for truncation.

Rules for Hostname Truncation

1. Delete any invalid characters (including spaces) from the strings. Valid characters include the following *only*:
   - numbers

- letters (upper or lowercase)
- the dash (that is, the hyphen [-]) character

2. The string does not contain an IP token value (“stringToKeep”). Perform a simple truncate. That is, delete all trailing characters beyond the 15/63 limit.
3. The string contain an IP token value (“stringToKeep”). Truncate the suffix in the string until the hostname value is in compliance.
4. The suffix has been fully truncated and the hostname length is still not in compliance. Truncate the prefix until the hostname value is in compliance, then replace the last character in the prefix with a dash.

Hostname Token Substitution and Truncation Examples

The following table shows the result of Cloud Manager retaining the “string to keep” and truncating to a normal, compliant hostname length.

Token	Token Substitution and Normalization
`%BS_Name%-%Short_IP%`	`LAMP-127`
`%Org_Name%-%Long_IP%`	**Linux:** `DigitalAirlines-67-122`
	Windows: `DigitalA-67-122`

14.5 Monitoring Resource Capacity

Roles that Can Perform This Task: Cloud Administrator, Zone Administrator (zone only), Approver (organization only)

The Capacity view provides information about used, reserved, and allocated resource capacity for organizations and zones.

- Section 14.5.1, “Accessing the Capacity View,” on page 155
- Section 14.5.2, “Understanding the Capacity View,” on page 155
- Section 14.5.3, “Updating the Capacity Data,” on page 158
- Section 14.5.4, “Debugging Capacity Collection Issues,” on page 158

14.5.1 Accessing the Capacity View

To open the Capacity view:

1 On the main navigation bar, click *Capacity*.

14.5.2 Understanding the Capacity View

The Capacity view includes three main sections:Editi

- “Capacity Summary Bar” on page 155
- “Organizations or Zones List” on page 156
- “Organizations or Zones Details” on page 157

Capacity Summary Bar

The Capacity Summary bar provides a summary for the total resources within your management scope.



For example, if you are a Cloud Administrator, you see a summary for all of the resources in your Cloud. If you are a Zone Administrator, you see a summary of all of the resources in the zones you manage. If you are an Approver, you see a summary of all of the resources in your assigned organizations and zones.

Each resource (Memory, CPU, and Storage) has its own capacity indicator. The indicator displays the used and reserved capacity as a percentage of the total available capacity.

The color of the indicator is determined by the Warning and Problem thresholds set for the Cloud environment. Green indicates that no thresholds have been reached, yellow indicates that the Warning threshold has been reached, and red indicates that the Problem threshold has been reached.

The Issues section of the Capacity Summary bar shows the following:

- No resource issues.
- The number of resources that have reached the Warning threshold.
- The number of resources that have reached the Problem threshold.

Organizations or Zones List

The Organizations or Zones list displays the organizations or zones for which you can view resource capacity.



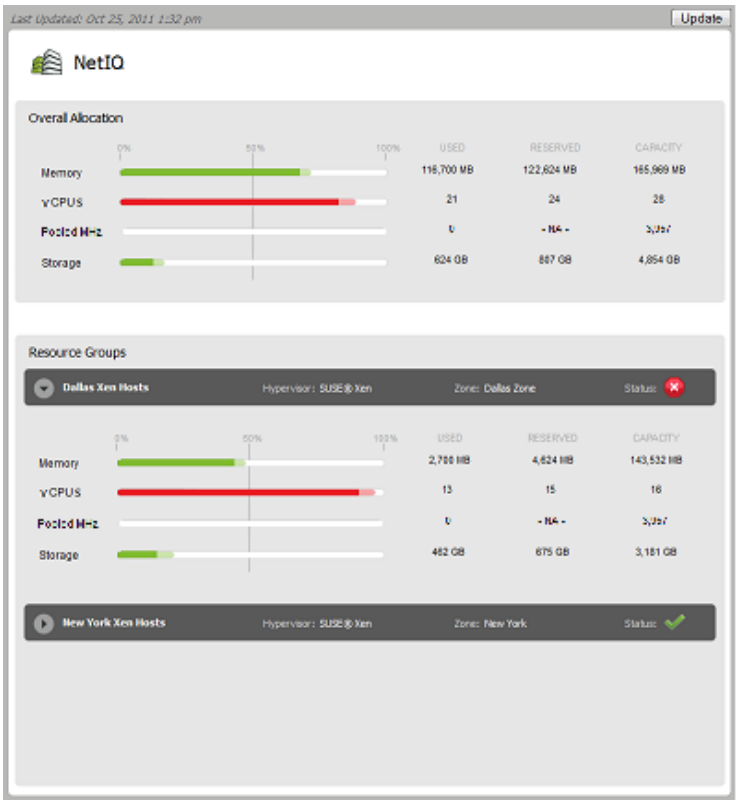
Depending on your role, you might be able to see either organizations or zones but not both.

The icon next to each organization or zone indicates the current status of the resources for that item. The statuses correspond to the statuses that can be listed in the Capacity Summary bar:

- No resource issues.
- One or more resources for the organization or zone have reached the Warning threshold.
- One or more resources for the organization or zone have reached the Problem threshold.

Organizations or Zones Details

The Organizations or Zones Details panel displays the resource capacity information for the organization or zone that is selected in the list.



The *Overall Allocation* section provides a summary of the used, reserved, and available capacity for the entire organization or zone. The *Resource Groups* section shows the same type of information for the each resource group in the organization or zone.

- **Memory:** The memory (RAM) resources.
- **vCPUs:** The host and cluster virtual CPU resources, represented in number of CPUs. This resource is displayed only if a resource group contains hosts or clusters.
- **Pooled MHz:** The resource pool CPU resources, represented in MHz of CPU. This resource is displayed only if a resource group contains vSphere resource pools.
- **Storage:** The shared storage resources. This includes SAN (Storage Area Network) and NAS (Network-attached Storage). Local storage is not included.

For each resource, the following information is displayed:

- **Used:** The amount of the resource that is actually being consumed by deployed workloads. For example, a workload might be allocated 4 GB of memory but only be using 2 GB.
- **Reserved:** The amount of the resource that is reserved for deployed workloads. For example, if a workload is allocated 4 GB of memory, all 4 GB are reserved.
- **Capacity:** The total amount of the resource that is available for deployed workloads. For memory from hosts and clusters, the *Capacity* field reflects the total physical memory of the hosts, not just the memory that is available for workloads. For example, a host's physical memory might be 16252 MB, with the actual memory available for workloads being 10695 MB.

Because the *Capacity* field uses the physical memory amount (16252 MB), you might not have as much memory capacity as appears. For memory from resource pools, the *Capacity* field reflects the actual memory pool size.

The *Resource Groups* sections provide capacity details for each resource group assigned to the organization or zone. The status of each resource group is also displayed.

The *Last Updated* field displays the last time the capacity data was updated. The *Update* button lets you update the data.

14.5.3 Updating the Capacity Data

The Capacity engine collects capacity data on a regular interval specified in System Configuration. The collected data is cached on the Cloud Manager Application Server.

The Capacity data is static, meaning that the Capacity view displays the same data until you update from the cached data on the Cloud Manager Application Server.

To update the data:

1. Click *Update*.

   If the cached data is newer than the current data, the Capacity view is updated with the cached data.
2. Click *Yes* to confirm that you to continue with the manual update.

   The Capacity engine collects the new data from the system. As soon as the data is collected, the Capacity view automatically updates to the new data.

14.5.4 Debugging Capacity Collection Issues

As a Cloud Manager administrator, you might occasionally encounter capacity collection issues in the Cloud Manager system. You can enable the debugging for capacity collection logging on your production level Cloud Manager Application Server to help you or NetIQ Support troubleshoot these issues.

The Cloud Manager Application Server uses a custom properties file, `/etc/cmplanner.logging.properties`, to enable logging. Its contents, including the `DEBUG` setting, look like this:

```
# Set root logger level to DEBUG and its only appender to A1.
log4j.rootLogger=WARN, CMFILE, CMOUT

# Planner is set to be a ConsoleAppender.
log4j.appender.CMOUT=org.apache.log4j.ConsoleAppender
log4j.appender.CMOUT.layout=org.apache.log4j.PatternLayout
log4j.appender.CMOUT.layout.ConversionPattern=%-4r [%t] %-5p %c %x - %m%n

# File appender
log4j.appender.CMFILE=org.apache.log4j.RollingFileAppender
log4j.appender.CMFILE.layout=org.apache.log4j.PatternLayout
log4j.appender.CMFILEt.layout.ConversionPattern=[%d{dd MMM yyyy HH:MM:ss}] %-5.5p
| %-16.16t | %-32.32c{1} | %X{bundle.id} - %X{bundle.name} - %X{bundle.version} |
%m%n
log4j.appender.CMFILE.file=${karaf.base}/logs/cloudmanager_server.log
log4j.appender.CMFILE.append=true
log4j.appender.CMFILE.maxFileSize=10MB
log4j.appender.CMFILE.maxBackupIndex=10


# Logging category modifications for novell-esb
log4j.logger.org.quartz=DEBUG
log4j.logger.com.netiq.rest.planner.CmPlanner=DEBUG
log4j.logger.com.netiq.service.planner.capacitybot.CapacityBot=DEBUG
```

To enable debug logging,

1. Edit the properties file, changing any `WARN` value to a `DEBUG` value (see example, above).
2. In the Karaf console, run an update on the AppServices bundle.

   The update process requires that you enter a module ID. You can find this ID using a `list` command and then browsing the list of modules and their accompanying IDs.

15 Managing Business Services

The following sections provide information to help you deploy business services, manage deployed business services, and change deployed business services:

- Section 15.1, “Requesting Business Services,” on page 161
- Section 15.2, “Managing Business Service Requests,” on page 163
- Section 15.3, “Importing a Single Virtual Machine into an Existing or a New Business Service,” on page 165
- Section 15.4, “Importing Multiple Virtual Machines into Cloud Manager Business Services,” on page 166
- Section 15.5, “Completing Pre-Build Configuration Tasks,” on page 175
- Section 15.6, “Completing Post-Build Configuration Tasks,” on page 175
- Section 15.7, “Managing Deployed Business Services,” on page 176
- Section 15.8, “Changing Deployed Business Services,” on page 179
- Section 15.9, “Extending Business Service Expiration Dates,” on page 182
- Section 15.10, “Displaying or Hiding Business Service Costs,” on page 183
- Section 15.11, “Understanding Business Service Contracts,” on page 184

15.1 Requesting Business Services

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

A deployed business service starts as a request. The request defines the business service, including its name, contract period, and workloads. As soon as you submit the request, it goes through an approval and build process that you can track until the business service is deployed.

1. On the main navigation bar, click *Business Services*, then click the *Requests* tab.
2. Click *Create* to display the Create Business Service Request dialog box.
3. Provide the following details for the business service:

   Service Name: Specify a name that is different than any other business service names for the business group. You can use letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, and period. The name must begin with a letter or number and have a maximum length of 110 characters.

   Start Date: Click to select the date you want the business service to be available.

   Expiration Date: Click to select the date you want the business service to no longer be available. If you don’t want the business service to expire, delete the date from the field.

   Contract Length: If an expiration date is selected, this field displays the total number of months for the contract.

   Organization: Select the organization for which you are creating the business service. You can select the business group before the organization, in which case the correct organization is automatically selected.

Business Group: Select the business group for which you are creating the business service.

Creator: Displays your user name.

Business Purpose: Provide details that explain the purpose or justification for the business service. This information is visible to the business service's approvers during the request approval process.

4 Add a workload to the business service:

4a Click *Add* to display the Select Workload Template dialog box.

The dialog box displays a list of workload templates. A template provides the base settings and costs for the workload.

4b Select the workload template from which to create the workload, then click *OK* to display the Configure Workload dialog box.

4c On the *Resources* tab, specify a name for the workload, select a service level, and customize the resources allocated to the workload.

You cannot customize the resources if they are locked. For additional information about the *Resources* settings, click the button.

4d Click the *Disks* tab to configure the workload's disks.

The workload's mandatory disks are listed. The mandatory disks are created with the workload and cannot be customized.

If the *Add* action, located above the list, is available, there are optional disks you can create for the workload. Click *Add*, then specify the size for the disk.

For additional information about the *Disks* settings, click the button.

4e Click the *Networks* tab (if it is displayed), select a network interface card, then assign the network and configure the network address and name servers.

For additional information about the *Networks* settings, click the button.

4f (Optional) Click the *Windows Settings* tab (if it is displayed), then provide an Administrator account password, computer name, and workgroup or domain information.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads. For additional information about the *Windows Settings* settings, click the button.

4g (Optional) Click the *Windows Licensing* tab (if it is displayed), then provide a Windows product key, and registration name.

If the Cloud Administrator has pre-populated the Windows Product Key field, the data is masked. You cannot copy this data.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide during pre-build configuration of the workloads. For additional information about the *Windows Licensing* settings, click the button.

4h (Optional) Click the *Linux Settings* tab (if it is displayed), then provide a host name, a domain name, and a Root password.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads.

For additional information about the *Linux Settings* settings, click the button.

4i (Optional) Click the *Console Access* tab, then set the password for VNC access to the workload.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads. For additional information about the *Console Access* settings, click the button.

4j Click *OK* to save the workload and add it to the *Workloads* list.

5 Add any additional workloads to the business service.

You can add additional workloads by repeating Step 4 or by copying an existing workload and then modifying it as necessary. To copy a workload:

5a Select the workload to copy, then click *Copy*.

5b Select the number of copies to create, and provide a unique name for each copy.

5c Click *OK*.

The new workloads are added to the *Workloads* list.

5d Edit each new workload to provide any missing information.

Each new workload contains as much of the original information as possible, but information such as network addresses, Windows computer names, and Linux host names are not copied because they need to be unique for each workload.

6 (Optional) Give other users ownership rights to the business service.

The Users list lets you see any users who have been explicitly assigned ownership rights to the business service. It does not show users who inherit ownership rights to the business service through their roles.

6a Click the *Users* tab, then click *Add*.

6b Select users from the two lists.

The *Members* list displays all users who are members of the business service's organization. The *System Users* list displays users who are not members of the organization.

You can Shift-click and Ctrl-click to select multiple users (or user groups).

6c Click *OK*.

The users are added to the list.

7 Click *Save* to add the request to the *Requested Services* list without submitting it.

or

Click *Submit* to add the request and submit it for approval.

15.2 Managing Business Service Requests

After you request a business service, a business service request is added to your *Requests* list. You can complete any of the following tasks to manage the request.

- Section 15.2.1, "Submitting a Request," on page 164
- Section 15.2.2, "Editing a Request," on page 164
- Section 15.2.3, "Withdrawing a Request," on page 164
- Section 15.2.4, "Deleting a Request," on page 165

15.2.1 Submitting a Request

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

The following steps assume that you have a saved business service request that you want to submit. If you have not yet created the request, see Requesting Business Services.

1. On the main navigation bar, click *Business Services*, then click *Requests*.
2. Select the request, then click *Submit*.

15.2.2 Editing a Request

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

You can edit only unsubmitted requests. If you want to change a submitted request, you must withdraw it, change it, and then resubmit it. For information about withdrawing a request, see Withdrawing a Request.

1. On the main navigation bar, click *Business Services*, then click the *Requests* tab.
2. Select the request, then click *Edit*.

   or

   Double-click the request.
3. Modify the request as desired.

   For a description of the settings, click the button.
4. Click *Save* to save your changes without submitting the request.

   or

   Click *Submit* to save your changes and submit the request.

15.2.3 Withdrawing a Request

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

Withdrawing a request stops the provisioning process. You can withdraw a submitted request until it reaches a certain phase in the *Building* process. At that point, the *Withdraw* action is no longer available.

After you withdraw a request, you can make changes to it and resubmit it, or you can delete it.

1. On the main navigation bar, click *Business Services*, then click the *Requests* tab.
2. Select the request, then click *Withdraw*.

15.2.4 Deleting a Request

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

You can delete non-submitted requests. To delete a submitted request, you must first withdraw it so that it becomes a non-submitted request. For information about withdrawing a request, see Withdrawing a Request.

1. On the main navigation bar, click *Business Services*, then click the *Requests* tab.
2. Select the request, then click *Delete*.
3. Click *Yes* to confirm the deletion.

15.3 Importing a Single Virtual Machine into an Existing or a New Business Service

Roles that Can Perform This Task: Cloud Administrator

If you have an existing virtual machine that is not part of the Cloud Manager system, you can import it as new business service or you can import it into a deployed business service. Because the virtual machine is already provisioned, the request and provisioning processes are skipped and the virtual machine is imported directly as a deployed workload.

In order to import a virtual machine, it must reside on a host associated with a resource group. The virtual machine can be running, shut down, or suspended.

The following steps explain how to import a virtual machine to become a workload inside a Cloud Manager business service. To import multiple virtual machines, see Section 15.4, “Importing Multiple Virtual Machines into Cloud Manager Business Services,” on page 166.

1. On the main navigation bar, click *Business Services*, then click *Virtual Machines*.
2. Click *Import* to display the Import Business Service dialog box.
3. Provide the following details for the business service:

   Service Name: Specify a name for the business service. The name should be unique among other business services you have created.

   Creator: Displays your name. You cannot change the creator.

   Created Date: Defaults to the current date. You cannot change the creation date.

   Expiration Date: Click to select the date you want the business service to no longer be available. If you don’t want the business service to expire, leave the field empty.

   If you select an expiration date, be aware that the business service is then considered “under contract” until the expiration date is reached. A business service under contract will not change price during the contract. In addition, you cannot make certain changes to business services that are under contract. For more information, see Section 15.11, “Understanding Business Service Contracts,” on page 184.

   Organization: Select the organization for which you are importing the business service.

   Business Group: Select the business group for which you are importing the business service.

Business Purpose: Provide details that explain the purpose or justification for the business service.

4 Import and configure a virtual machine:

4a On the *Workloads* tab, click *Import* to display the Select Virtual Machine dialog box.

The dialog box displays a list of virtual machines that can be individually imported into the Cloud Manager environment and associated to selected business group.

4b Select the virtual machine to import, then click *OK* to display the Configure Imported Workload dialog box.

4c Configure the workload you intend to import:

4c1 On the *Resources* tab, configure the following:

Workload Name: By default, the workload inherits the virtual machine name, but the name is editable. Make sure the name is different from other workloads included in the business service.

Workload Template Name: By default, the workload template is derived from the virtual machine template name in the Orchestration Server. Open the drop-down and select the VM template name from the list. If there are no options listed, no VM template was created for it in the Orchestration Server prior to the import.

License Costs: If available (usually because no workload template is associated with the workload), specify the monthly cost for licenses associated with the workload.

Service Level: Select the default service level for the business service.

Package: If available (usually because the workload is associated to a workload template), select a resource pricing package for the business service.

4c2 (Optional) On the *Disks* tab, select a disk, then enable the *Cost* check box to include the disk's cost in the business service cost or disable the check box to exclude the disk from the business service cost.

NOTE: Because the virtual machine was previously configured in the hypervisor environment where it was created, no other settings on any other tabs are configurable. You can change the settings of the VM when it is imported to become a workload inside Cloud Manager.

4c3 Click *OK* to save the configuration and add the workload to the business service.

5 Click *Import* to add the business service to the list of deployed business services.

After the import, you can change the workload's resources if necessary. This is a change to the business service and requires a change request to go through the approval process and the imported workload to be rebuilt. See Chapter 15.8, "Changing Deployed Business Services," on page 179.

15.4 Importing Multiple Virtual Machines into Cloud Manager Business Services

---

Roles that Can Perform This Task: Cloud Administrator

The following methods explain how to import multiple virtual machines to become workloads inside a number of Cloud Manager business services.

- Section 15.4.1, “Bulk Importing Multiple Virtual Machines,” on page 167
- Section 15.4.2, “About the Bulk Import Spreadsheet,” on page 168

To import a single virtual machine, see Section 15.3, “Importing a Single Virtual Machine into an Existing or a New Business Service,” on page 165.

15.4.1 Bulk Importing Multiple Virtual Machines

Because your data center environment probably has one or more hypervisors where many existing virtual machines are already configured and running, you should take full advantage of Cloud Manager’s Orchestration components to discover these machines and list them as unassigned virtual machines. These unassigned VMs have the potential of being fully imported into the Cloud Manager environment, where they can be assigned as workloads and associated to Cloud Manager business services, business groups, and organizations.

Use these steps to import multiple, unassigned VMs discovered by the Cloud Manager Orchestration components:

1. On the Cloud Manager Web Console main navigation bar, click *Business Services*, then click *Virtual Machines* to display the *Unassigned Virtual Machines* list.
2. From the list, select the virtual machines that you want to import (press and hold the Ctrl key to multi-select).
3. Click *Export* to move the existing virtual machine configuration data to a Microsoft Excel (`.xls`) spreadsheet.
4. Open the exported spreadsheet with a compatible spreadsheet application, then for each VM (that is, each row in the spreadsheet), enter additional data. At a minimum, you must enter information into the gray-shaded columns; these are required.

   For more information about the export VM spreadsheet, see “About the Bulk Import Spreadsheet” on page 168.
5. Validate the virtual machine configurations:
   - **5a** On the *Virtual Machines* tab, click *Bulk Import* to display the Bulk Import dialog box.
   - **5b** Browse to and download the bulk import spreadsheet file you exported and edited in Step 3 and Step 4.
   - **5c** Click *Validate* to upload the data to the Application Server, where it is processed and evaluated for validity.
   - **5d** Check the validation results either in the Bulk Import dialog box or in the spreadsheet, then correct the errors in the spreadsheet and save it again.
   - **5e** Repeat all of the steps Step 5b through Step 5d until every row in the spreadsheet is valid.
6. Click *Import* to add the virtual machines to the Cloud Manager environment (that is, as workloads associated to business services, business groups, and organizations).

TIP: If you have incorporated NetIQ Access Manager authentication into your Cloud Manager environment, a two minute HTTP timeout is enforced on your browser (that is, the Cloud Manager Web console where you observe the import). When the limit is reached, the Bulk Import dialog box displays a `Ready to Import` message, which you can ignore: the process actually continues on the Application Server until completion. Close the Bulk Import dialog box instead. (In this scenario, if you click *Get Results* instead of closing the dialog box, the server returns inaccurate import data. If you click *Import*, the server generates specious import errors.)

After you close the dialog box, wait for a server-generated email message with an attached `.pdf` file. The email signals the end of the import process and the `.pdf` details its results.

If you want to change the Access Manager HTTP timeout, specify a different value for *Data Read Timeout* field, as explained in the *Configuring TCP Listen Options for Clients* section of the *NetIQ Access Manager 3.2 SP2 Access Gateway Guide*.

7 Click *Close* to close the Bulk Import dialog box.

At this point, you can change the imported workload's resources if necessary. Making such changes constitutes a change to the business service(s), and so this requires one or more change requests to go through the approval process and the imported workload(s) to be rebuilt. See Chapter 15.8, "Changing Deployed Business Services," on page 179.

Using the Get Results Function

After you browse to and perform an initial validation of the bulk import spreadsheet, the *Get Results* function becomes available on the Bulk Import dialog box. This function can be accessed at any time. It also creates an `.xls` spreadsheet: one that provides more exact results of the current state of the VM import process, including:

- The deletion of superfluous rows (that is, either blank or containing information that is not used in the validation)
- More verbose results of the validation status.
- Workload IDs assigned to the VMs after an initial import.

Get Results is useful if you experienced a successful import process and you want to keep a detailed record of that event.

Ongoing Bulk Import Issues

When an administrator performs a bulk import of VMs, one column in the bulk import spreadsheet allows the inclusion of an add-on service with the VM, giving the impression that information for the add-on service will be put through to the workload after import. Although a VM be imported successfully with this designated add-on service, Cloud Manager does not recognize its add-on configuration and throws an error:

```
Additional Information is needed.
```

This indicates that no AOS information was actually included during the import. You will need to edit the imported workload and add text for the user that instructs him or her to provide configuration instructions to the build administrator. For more information, see Section 12.6, "Add-On Services," on page 121.

15.4.2 About the Bulk Import Spreadsheet

Many large data centers keep detailed spreadsheet or .csv records of their virtual machines and the details of the respective VM configurations. If your enterprise keeps such records, you or someone on your staff can map that VM data to the Cloud Manager VM export spreadsheet. Cloud Manager uses this specially formatted .xls spreadsheet as the principal mapping tool to import unassigned VMs to the Cloud Manager environment, where they are configured by the Cloud Manager Application Server to become Cloud Manager workloads.

If you want to perform such a mapping, perhaps from `.csv` format to `.xls` format, we suggest that you initially choose one VM from the *Unassigned Virtual Machines* list, create an export file (see Step 3 above), perform the mapping based on your analysis of the Cloud Manager spreadsheet, then

delete the populated data for the single VM before you import your mapped `.csv` file with data from all the VMs you want to import. If auto-generating a .csv file destined to be converted into a .xls file, many of the column headings are examined to verify the spreadsheet's structure. Be sure to match column headings exactly to ensure a successful import into Cloud Manager.

You can also choose to export many or all of the Cloud Manager unassigned VMs to the spreadsheet and then manually adjust the configuration of each virtual machine according to the requirements for import.

When first launched, many of the columns of the spreadsheet are prepopulated with data discovered by the Orchestration Server and available in the respective VM details. Do not change the column order or the prepopulated columns. The column names are helpful for you to remember the purpose of the respective VM data. Other things you should know about the spreadsheet include the following:

- Many of the column heads include associated comments that help you understand the purpose of the data in the column and whether it is required as part of the mapping process. You can identify those columns by the small red rectangle in the upper right-hand corner of the cell and by its gray shading. Hover over a column to display the comment. Look for the word *required* in the comment header.
- When you complete your changes (which can be accomplished manually or through automatic merging of other spreadsheet data) to the individual columns, use the *Validate* function to upload the data to the Application Server. The server validates the contents, adds missing details, and provides validation messages inside the spreadsheet. You can keep the spreadsheet open and make changes while using the Validate function repeatedly.
- The bulk import spreadsheet is also a worksheet. The *Validation Status* column of the spreadsheet provides a list of reasons preventing the VM from being imported. Use this information to adjust the configuration for the VMs that show errors, then click *Get Results* to revalidate until all errors are corrected.
- If you don't know database IDs, provide the name, you might want names for other people looking at the spreadsheet
- IDs are fast lookup vs. names
- provide IDs ahead of time if you can / providing both is good for both human and machine (lookup)

Required Columns in the Spreadsheet

The bulk import spreadsheet includes many columns, only a few of which are required. The data you provide in the required columns is passed to the Application Server, where the actual VM to workload configuration and conversion takes place. The information you provide in the non-required columns is informational only, and although passed to the workload, it is not essential to its configuration.

NOTE: The exception to the non-essential data is column 3: *Validation Status*. Although the information in this column is not required for configuring the workload, it is provided by the Application Server, and is essential for your guidance during the validation process as you adjust the column information for each row (that is, VM) listed in the spreadsheet.

The following table provides more details about the columns of a typical bulk import spreadsheet:

Column Name	Represents	Required	Details About the Cell
VM Name or VM ID	Virtual Machine name	Yes	◆ If you use the VM Export tool to create the spreadsheet, Cloud Manager populates the cell for you. ◆ If you create your own spreadsheet, you must enter the name exactly as it appears in the Orchestration Server zone. The lookup function of the import process is case sensitive.
Validation Status	Bulk Import validation status	Provided by the Cloud Manager Application Server	◆ Provides the list of reasons preventing the current virtual machine from being importable. ◆ To see the validation status, use the *Get Results* button on the bulk import dialog box.
Zone Name or Zone ID	Name of the Cloud Manager Orchestration Zone containing the VM	Yes	◆ Case sensitive. ◆ Enter the name of the CMOS zone containing the VM. ◆ Provided when exporting VMs from Cloud Manager.
Org Name or Org ID	Cloud Manager Organization Name	No (Conditional)	- Case sensitive. - This field is required. However, it can be derived from the provided pre-existing Business Service or Business Group. If it can be derived from these other objects, you do not need to enter the value here.
BG Name or BG ID	Cloud Manager Business Group Name	No (Conditional)	- Case sensitive. - This field is required. However, it can be derived from the provided pre-existing Business Service. In that case, it is not required.
BS Name or BS ID	Cloud Manager Business Service Name	Yes	- Case sensitive if used to find a pre-existing Business Service. Not case sensitive if used to create a new Business Service. - If creating a new Business Service, the first spreadsheet row referencing the BS will create the Business Service object. Subsequent rows referencing the same BS will import into the previously created BS.
BS Purpose	Cloud Manager Business Service Purpose	No	no comment
License Cost	Cost to License the Workload	No	no comment
BS Expire Date	Cloud Manager Business Service Expiration Date	No	Date Format Date must be entered in m/d/yyyy format

Column Name	Represents	Required	Details About the Cell
BS Custom 1[(1)]	Custom Field Added to the Cloud Manager Business Service	No	no comment
WL Name	Cloud Manager Workload Name	Yes	- Enter the desired name of the created Workload object. - Default to the VM Name when exporting VMs from Cloud Manager
WL ID	Cloud Manager Workload ID	Provided	Provided after import. The created workload ID will be written to the results spreadsheet when successfully imported. If providing the Get Results spreadsheet as the input to a subsequent Import, any rows with an existing, valid Workload ID will be skipped without errors.
WLT Name or WLT ID	Cloud Manager Workload Template Name	No	- Case sensitive. - Must match VM details (os type, hypervisor, & resource ranges). - If multiple RPPs are associated with the WLT, the RPP column is then required.
RG Name or *RG ID*	Cloud Manager Resource Group Name	No (Conditional)	- Case sensitive. - Specified RG must be available to Business Group. - RG's hypervisor must match virtual machine's hypervisor - Required if multiple resource groups match the VM's details. - If not provided, Cloud Manager will attempt to discover the matching Resource Group, and will succeed if only 1 match is found.
RPP Name or *RPP ID*	Cloud Manager Resource Pricing Package Name	No (Conditional)	- Case sensitive. - Must be available to the Workload Template identified in the WLT columns - Must match the resource ranges of the WLT and the actual resource allocations of the VM.
SL Name or *SL ID*	Name of the Cloud Manager Service Level to Be Associated to this Workload	Yes	- Case sensitive. - Must be available to the Resource Group identified in the RG columns - Required if multiple Service Levels match the Resource Group. - If not provided, Cloud Manager will attempt to discover the matching Service Level, and will succeed if only 1 match is found.

Column Name	Represents	Required	Details About the Cell
VNC Password	Password used to gain remote access to this VM's VNC Console	No	- Case sensitive - Passwords are never written to the spreadsheet when exported. - Providing a password during import may cause the VM to be restarted (Confirm with Norm)
Computer Name	Computer Name for Windows systems Hostname for Linux systems	No	- For Windows VMs, provide the computer name. - For Linux VMs, provide the hostname
Domain Name	Windows Domain Name	No (Conditional)	- Recommended for Linux VMs - Required for domain-based Windows VMs - For workgroup-based Windows VMs, leave blank and enter workgroup name in next column
Win Workgroup	Windows Workgroup Name	No (Conditional)	- Required for workgroup-based Windows VMs - Ignored for Linux systems
Win Admin Password	Windows Administrator Password	No	- Case sensitive. - Ignored for Linux systems
Win Domain User ID	Windows Domain Administrator User ID	No	- get from help system - Ignored for Linux systems
Win Domain Password	Windows Domain Administrator password	No	- get from help system - Ignored for Linux systems
Win Product Key	Windows Product Key	No	- get from help system - Ignored for Linux systems
Win RegisteredTo Name	Windows Registered To Name	No	- get from help system - Ignored for Linux systems
Run Once Commands	Windows Run Once Commands	No	- get from help system - Ignored for Linux systems

Column Name	Represents	Required	Details About the Cell
NIC1 IP DHCP[(2)]	Use DHCP for IP Address on NIC 1	No	- Boolean value must be either TRUE or FALSE - If not provided in the spreadsheet during import, puts the IP DHCP setting for this NIC into an indeterminate state. Later, when changing network settings during a change request, users will be required to specify the IP DHCP setting. - Create additional NIC# sections as appropriate.
NIC1 IP[(2)]	IP Address for NIC 1	No	- Enter the IP Address for this NIC - Provided when exporting VMs from Cloud Manager
NIC1 MAC[(2)]	MAC Address for NIC 1	No (Conditional)	- Enter the MAC Address for this NIC - Required if providing NIC details during import - Provided when exporting VMs from Cloud Manager
NIC1 Network Name[(2)] or NIC1 Network ID	Network associated with NIC 1	No	- Case sensitive - Must match a network in the specified CMOS Zone - Provided when exporting VMs from Cloud Manager
NIC1 Mask[(2)]	Subnet Mask for NIC 1	No	- Provided when exporting VMs from Cloud Manager
NIC1 DNS DHCP[(2)]	Use DHCP for DNS Servers and Suffixes for NIC 1	No	- Boolean value must be either TRUE or FALSE - If not provided in the spreadsheet during import, puts the DNS DHCP setting for this NIC into an indeterminate state. Later, when changing network settings during a change request, users will be required to specify the DNS DHCP setting.
NIC1 DNS Servers[(2)]	Domain Name Servers for NIC 1	No	- Enter the IP address of one or more DNS servers - Place the IP Address of each DNS Server on a separate line within the cell - Provided when exporting VMs from Cloud Manager

Column Name	Represents	Required	Details About the Cell
NIC1 DNS Suffixes[(2)]	DNS Suffixes for NIC 1	No	- Enter one or more DNS suffixes - Place each suffix on a separate line within the cell - Provided when exporting VMs from Cloud Manager
NIC1 Gateways[(2)]	Gateways for NIC 1	No	- Enter the IP Address of one or more gateways - Place each address on a separate line within the cell - Provided when exporting VMs from Cloud Manager
AddOn1 Name[(3)] or AddOn1 ID	Add-on Application or Service	No	- Case sensitive - Enter the name or ID of an Add-on Application or Service to be associated to this Workload in Cloud Manager - Must be available to the Workload Template specified in the WLT columns - Note that “locked” Add-ons are automatically associated and do not need to be specified here. - Create additional AddOn# sections as appropriate.

(1): You can add customized fields to every Cloud Manager business service. The fields can contain any information you want. You can add only three custom fields to the business service The first custom field is designated as BS Custom 1, the second as BS Custom 2, and the third as BS Custom 3.

(2): A workload can have multiple network cards, each of which has a unique network configuration. The first network card is designated as NIC1, the second as NIC2, and so on. The spreadsheet must provide a complete set of 10 NIC columns for each NIC. For convenience, these are color coded during the export process. Enough NIC “sections” must be defined for the greatest number of NICs in the importing list of VMs. The first NIC is designated as NIC1, the second as NIC2, and so on. Ensure column headings are labelled correctly.

(3): A workload can have multiple add-on applications or add-on services associated with it. Each can have a unique name and ID. Enough AddOn “sections” must be defined for the greatest number of AddOns in the importing VMs. The first add-on is designated as AddOn1, the second as AddOn2, and so on. Ensure column headings are labelled correctly.

15.5 Completing Pre-Build Configuration Tasks

After a business service request receives both Administrator and Sponsor approval, a *Pre-build configuration* task is sent to you and any other Build Administrators assigned to the business service's organization or business group. Cloud Administrators also receive the task.

The task pauses the provisioning workflow to give you an opportunity to perform any pre-build configuration tasks for the workloads. This might be a task that you complete in the Cloud Manager console, such as configuring the Windows settings (Administrator account password, computer name, product license key, and registered name) or remote console password for one of the workloads. Or, it might be a task that you need to perform at the Cloud Manager Orchestration Server or in your hypervisor tools.

1. On the main navigation bar, click *Tasks*, then click the *Unclaimed* tab.
2. Select the task (the subject is *Pre-build configuration of a new Business Service*)

   If there are configuration tasks that must be completed for the business service in the Cloud Manager console, the *Complete* action is not available and an asterisk (*) appears next to it. You can mouse over the asterisk to view the remaining tasks.
3. If you want to claim the task so that others cannot work on it, click *Claim*.

   Claimed tasks are moved to your *My Tasks* list.
4. Complete any pre-build configuration tasks required for the workloads:
   - **Cloud Manager console tasks:** Select the configuration task, then click the *Review* action to display the task. Select a workload that needs to be configured, then click *Edit*. Configure the required settings and save the workload.
   - **External tasks:** Complete the tasks.
5. Mark the task as complete:
   - If you claimed the task, it is in your *My Tasks* list. Select the task, then click *Complete*.
   - If you did not claim the task, it is in the *Unclaimed* list. Select the task, then click *Complete*.

15.6 Completing Post-Build Configuration Tasks

After a business service is successfully built, a *Post-build configuration* task is sent to you and any other Build Administrators assigned to the business service's organization or business group. The task is also sent to Cloud Administrators.

The task pauses the provisioning workflow to give you an opportunity to perform any post-build configuration tasks for the workloads. You must complete the configuration (if any) and then mark the task as complete to continue the provisioning workflow.

1. On the main navigation bar, click *Tasks*, then click the *Unclaimed* tab.
2. If you want to claim the task so that no others can work on it, select the task, then click *Claim*.

   Claimed tasks are moved to your *My Tasks* list.
3. If you want to review the business service request, select the task (the subject is *Post-build configuration of a new Business Service*), then click *Review*.
4. Complete any post-build configuration tasks required for the workloads.
5. Mark the task as complete:
   - If you claimed the task, it is in your *My Tasks* list. Select the task, then click *Complete*.
   - If you did not claim the task, it is in the *Unclaimed* list. Select the task, then click *Complete*.

15.7 Managing Deployed Business Services

You can cycle (start, suspend, stop) a deployed business service's workloads, open a remote console to a workload, and give other users access to the business service.

- Section 15.7.1, "Starting, Suspending, or Shutting Down a Workload," on page 176
- Section 15.7.2, "Opening a Workload Console," on page 176
- Section 15.7.3, "Adding Files to a Workload," on page 177
- Section 15.7.4, "Delegating Ownership of a Business Service," on page 178

15.7.1 Starting, Suspending, or Shutting Down a Workload

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service with the workload you want to cycle, then click *Manage*.
3. In the *Workloads* list, select the workload.
4. Click one of the following controls:

   Start the workload.

   Suspend the workload. This action takes a snapshot of the workload's disk and memory status, saving its current state.

   Shut down the workload.
5. Click *Close*.

15.7.2 Opening a Workload Console

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

Cloud Manager provides remote access to workloads via a VNC console.

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service with the workload you want to access, then click *Manage*.

   You can also access the workload from the *Workloads* tab.
3. In the *Workloads* list, select the workload, then click *Console*.
   - The Console action opens a new browser window that provides VNC access to the selected workload's desktop. The action is enabled even when a workload is shut down, which lets you enter a VNC password before the virtual machine begins to boot. You can therefore interact with the boot sequence very early in the process, potentially loading resources from a PXE server or similar network service.

- If the VNC session does not require a password, you can press Enter. If you enter a password incorrectly when the VNC session is established, you can enter the password again.
- The action is not enabled for workloads that have not yet been configured for VNC access. This is common for imported workloads.

4 Specify the console password to log in.

After you log in, Cloud Manager opens a remote console session to the workload. Most functions of the remote console work the same as the local console. However, if you need to use special key sequences such as Ctrl+Alt+Del, you must select them from the *Send Special Key Sequence* menu located at the top of the console.

NOTE: The VNC console is rendered using HTML 5. If you use Internet Explorer 9 as the browser to view the Cloud Manager Web Console, you need to make sure that either the Adobe (Shockwave) Flash add-on is enabled or that the Google Chrome Frame add-on is enabled. Both of these add-ons are available for free download.

You can toggle the console display between full-screen mode and windowed mode by using the F11 key.

5 When you have finished, close the browser window to end the session.

15.7.3 Adding Files to a Workload

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

---

Cloud Manager provides file uploads to workloads via a VNC console.

1 On the main navigation bar, click *Business Services*, then click *Deployed* > *Workloads*.

2 Select the business service with the workload you want to access, then click *Manage*.

3 In the *Workloads* list, select the workload, then click *Console*.

- The *Console* action opens a new browser window that provides VNC access to the selected workload's desktop. The action is enabled even when a workload is shut down, which lets you enter a VNC password before the virtual machine begins to boot. You can therefore interact with the boot sequence very early in the process, potentially loading resources from a PXE server or similar network service.
- If the VNC session does not require a password, you can press Enter. If you enter a password incorrectly when the VNC session is established, you can enter the password again.
- The action is not enabled for workloads that have not yet been configured for VNC access. This is a common situation for imported workloads.

4 Specify the console password to log in.

- After you log in, Cloud Manager opens a remote console session to the workload.
- You can toggle the console display between full-screen mode and windowed mode by using the F11 key.

5 From your local file system, select (or multi-select) the files you want to copy, then drag the selection to the browser window with the remote console session.

- The file or group of files mounts on the workload via the CD-ROM configured on the workload.

- Use the file manager on the workload to further copy the file or files from the CD-ROM drive to the location where you want them.
- "Eject" the CD using the icon at the top of the Remote Console window when you have finished managing the files.

6 When you have finished with the upload, close the browser window to end the session.

IMPORTANT: A user with access to the vSphere environment could eject the mounted CD while you are performing a file upload in Cloud Manager. If this should happen, click *Eject* on the VNC console of the workload where you are uploading.

For more details about how the file upload feature works, its limitations, and how to troubleshoot a file upload failure, see "The File Upload Feature" in the *NetIQ Cloud Manager 2.3 Component Reference*

File Upload Icons

The table below lists the icons that display on the workload console during the file upload procedure.

Icon	Description
	Remote console communications are secure.
	Remote console communications are *not* secure.
	The CD ROM is ready and working. Click this icon to "eject" the CD when you no longer need the files.
	The file upload is in process. Click this icon to stop the process.
	Click this icon for more information about the workload console.

15.7.4 Delegating Ownership of a Business Service

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

---

By default, the following users have ownership rights to a business service:

- The creator of the business service
- All Business Service Owners for the business group that the business service belongs to
- All Business Service Owners for the organization
- All Organization Managers

If necessary, you can delegate ownership to other users. When you do so, the user receives rights to view, manage, and change the business service.

1 On the main navigation bar, click  *Business Services*, then click *Deployed*.

2 Select the business service to which you want to give another user access, then click *Manage*.

3 Click the *Users* tab to display the *Business Service Owner* list.

4 Click *Add*, select the users or user groups you want to give access to the business service, then click *OK*.

You can Shift-click and Ctrl-click to select multiple users and user groups.

If the *Add* action is not available, you do not have rights to give access to other users.

15.8 Changing Deployed Business Services

You can change a business service's details (expiration date and business purpose) and workloads. Any change (including changing the expiration date and business purpose) generates a change request that must be approved before the changed business service can be redeployed.

- Section 15.8.1, "Changing a Business Service's Details," on page 179
- Section 15.8.2, "Reassigning a Business Service to a Different Business Group," on page 179
- Section 15.8.3, "Adding a Workload," on page 180
- Section 15.8.4, "Modifying a Workload," on page 181
- Section 15.8.5, "Removing a Workload," on page 182

15.8.1 Changing a Business Service's Details

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service, then click *Change*.
3. Modify the details for the business service:

   Expiration Date: Specify the date you want the business service to no longer be available.

   Business Purpose: Provide details that explain the purpose or justification for the change to the business service.
4. Click *Submit* to submit your business service change request.

   or

   If you don't want to submit the change request at this time, click *Save* to save the change request.

   A change icon is added next to the business service name in the *Deployed* list to indicate that a change request has been generated and added to the *Requests* list. If you saved the change request rather than submitting it, you must go to the *Requests* list to edit the request or submit it.

15.8.2 Reassigning a Business Service to a Different Business Group

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

You can reassign a business service to a different business group if 1) you have rights to the other business group and 2) the business group's resources (hypervisor, service levels, and so forth) support the business service's workloads.

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service, then click *Change*.
3. In the *Business Group* list, select the target business group.
4. Click *Submit* to submit your business service change request.

   or

   If you don't want to submit the change request at this time, click *Save* to save the change request.

   A change icon is added next to the business service name in the *Deployed* list to indicate that a change request has been generated and added to the *Requests* list. If you saved the change request rather than submitting it, you must go to the *Requests* list to edit the request or submit it.

15.8.3 Adding a Workload

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service to which you want to add a workload, then click *Change*.
3. On the *Workloads* tab, click *Add* to display the Select Workload Template dialog box.

   The dialog box displays a list of workload templates. A template is the basis for a workload, providing basic workload settings and costs.
4. Select the workload template from which to create the workload, then click *OK* to display the Configure Workload dialog box.
5. On the *Resources* tab, specify a name for the workload, select a service level, and customize the resources allocated to the workload.

   If the resources are locked, you cannot customize them. For additional information about the *Resources* settings, click the button.
6. Click the *Disks* tab to configure the workload's disks.

   The workload's mandatory disks are listed. The mandatory disks are created with the workload and cannot be customized.

   If the *Add* action, located above the list, is available, there are optional disks you can create for the workload. Click *Add*, then specify the size for the disk

   For additional information about the *Disks* settings, click the button.
7. Click the *Networks* tab (if it is displayed), select a network interface card, then assign the network and configure the network address and name servers.

   For additional information about the *Networks* settings, click the button.
8. (Optional) Click the *Windows Settings* tab (if it is displayed), then provide an Administrator account password, computer name, and workgroup or domain information.

   If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads.

   For additional information about the *Windows Settings* settings, click the button.

9 (Optional) Click the *Windows Licensing* tab (if it is displayed), then provide a Windows product key, and registration name.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads.

For additional information about the *Windows Licensing* settings, click the button.

If the Cloud Administrator has pre-populated the Windows Product Key field, the data is masked. You cannot copy this data.

10 (Optional) Click the *Linux Settings* tab (if it is displayed), then provide a host name, a domain name, and a Root password.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads.

For additional information about the *Linux Settings* settings, click the button.

11 (Optional) Click the *Console Access* tab, then set the password for VNC access to the workload.

If you do not provide this information at this time, a Build Administrator or Cloud Administrator must provide it during pre-build configuration of the workloads.

For additional information about the *Console Access* settings, click the button.

12 Click *OK* to save the workload and add it to the *Workloads* list.

13 Click *Submit* to submit your business service change request.

or

If you don't want to submit the change request at this time, click *Save* to save the change request.

A change icon is added next to the business service name in the *Deployed* list to indicate that a change request has been generated and added to the *Requests* list. If you saved the change request rather than submitting it, you must go to the *Requests* list to edit the request or submit it.

15.8.4 Modifying a Workload

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

1 On the main navigation bar, click *Business Services*, then click *Deployed*.

2 Select the business service with the workload you want to modify, then click *Change*.

3 On the *Workloads* tab, select the workload to modify, then click *Edit*.

4 Modify the workload settings as desired.

NOTE: The Linux Root Password is not editable on a change request.

For information about the workload settings, click the button.

IMPORTANT: The *Service Level* field and *Resource Group* field (if it is available) determine which resource group hosts the workload. If you use either of these fields to change the workload's resource group, you need to stop and then start the workload after it is rebuilt. The workload is not moved to the new resource group until it is manually restarted.

5 Click *OK* to save the workload changes.

6 Click *Submit* to submit your business service change request.

or

If you don’t want to submit the change request at this time, click *Save* to save the change request.

15.8.5 Removing a Workload

Roles that Can Perform This Task: Cloud Administrator, Organization Manager, Business Service Owner

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service with the workload you want to remove, then click *Change*.
3. On the *Workloads* tab, select the workload, then click *Remove*.
4. Click *Yes* to confirm the action.
5. Click *Submit* to submit your business service change request.

   or

   If you don’t want to submit the change request at this time, click *Save* to save the change request.

15.9 Extending Business Service Expiration Dates

Roles that Can Perform This Task: Cloud Manager, Organization Manager, Business Service Owner

To extend the expiration date for a business service, you must submit a business service change request. The request goes through the standard approval and configuration workflow for a business service.

1. On the main navigation bar, click *Business Services*, then click *Deployed*.
2. Select the business service, then click *Change*.
3. Change the expiration date to the desired date.
4. Click *Submit* to submit your business service change request.

   or

   If you don’t want to submit the change request at this time, click *Save* to save the change request.

   A change icon is added next to the business service name in the *Deployed* list to indicate that a change request has been generated and added to the *Requests* list. If you saved the change request rather than submitting it, you must go to the *Requests* list to edit the request or submit it.

15.10 Displaying or Hiding Business Service Costs

The Cloud Manager console displays business service costs in places such as the business service lists, the individual business service dialog boxes, and business service cost reports. System users always see the business service costs. However, you can choose whether or not to display costs to Organization members.

There are three levels at which you can configure the *Costs* setting:

- **Organization:** This setting applies to all Organization members. For example, if the *Costs* setting is set to *Hide*, no Organization members can see business service costs.
- **Business Group:** This setting overrides the organization setting for the business group. For example, if the organization *Costs* setting is set to *Hide*, but the business group *Costs* setting is set to *Show*, Organization members can see the costs for that business group's business services.
- **User:** Each user account has an *Always show costs* setting that overrides the organization and business group *Costs* setting.

15.10.1 Configuring an Organization's Costs Setting

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

1. On the main navigation bar, click *Organizations.*
2. On the *Organizations* tab, select the target organization, then click *Edit*.
3. In the upper-right corner of the dialog box, click the icon to display the Organization Configuration dialog box.
4. Under *Organizations Costs*, select *Show* or *Hide*.

   The setting applies to all Organization members unless it is overridden for a business group or individual user.
5. Click *OK* to save your changes.

15.10.2 Configuring a Business Group's Costs Setting

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

1. On the main navigation bar, click *Organizations.*
2. Click the *Business Groups* tab, select the target business group, then click *Edit*.
3. In the *Costs* field, select *Override* to override the setting inherited from the organization, then select *Show* or *Hide*.

   The setting applies to all business services created for the business group. If the costs are hidden, no Organization members can see the costs unless you enable the *Always show costs* setting for individual users.
4. Click *Save* to save your changes.

15.10.3 Configuring a User's Costs Setting

Roles that Can Perform This Task: Cloud Administrator, Organization Manager

1. On the main navigation bar, click *Users.*
2. Click the *Users* tab, select the target user, then click *Edit*.
3. In the *Cost Visibility* field, select *Always show costs* to enable the user to see business service costs regardless of the organization and business group *Costs* settings.

   The setting is disabled for System users. System users can always see business service costs.
4. Click *Save* to save your changes.

15.11 Understanding Business Service Contracts

Business services with an expiration date are considered "under contract" until the expiration date is reached. A business service under contract will not change price during the contract. In addition, you cannot make certain changes to business services that are under contract.

15.11.1 Restrictions Resulting from an Active Contract

Changes to a business service are restricted if the business service is under contract. The following restrictions apply when the business service is under contract:

- You can change the resources on the business service (that is, the number of CPUs, the amount of RAM, the number of NICs, and the total costable storage) only within the minimum/maximum limits defined on the associated package.
- You can change the associated resource pricing package only if it is an "upgrade," that is, the new package must require the minimum resources and/or contract length to be the same or greater value. A package "downgrade," that is, changing the settings to the minimum resources and/or contract length, is not allowed.
- You can only upgrade the service level of the business service to the same price or more expensive price. A downgrade of the service level is not allowed.
- You cannot reduce or delete an expiration date on a business service.
- You cannot delete a workload from the business service.
- You cannot delete a contracted business service.

15.11.2 Contract Change Period

When the business service nears its expiration date, you will receive an email notifying you that changes to the business service are no longer restricted. During this "change period" you can make any change, including deleting the business service.

15.11.3 Contract Renewals

Some resource pricing packages have minimum contract lengths. When you choose one of these packages for a workload, the business service must have a contract that is at least as long as the workload's minimum contract length. Some packages also include a contract renewal length. If a

business service contains a workload with a contract renewal length, the business service's contract renews for the specified length when it reaches the expiration date. In other words, when the business service's expiration date is reached, the expiration date is extended by the renewal length.

You can expect the following behaviors for contract renewals:

- If a business service has several workloads, the contract renewal length for the entire business service is derived from the workload with the longest contract renewal length.
- If a business service has no number of days for contract renewal, the business service expires.
- If a business service contract is renewed, the expiration date is extended. No cost changes occur.
- If a business service undergoes a change request during the change period, the expiration date is extended.
- If a business service change request is pending during a contract change period, its contract is not renewed until the change request is reviewed and approved.

16 Managing Tasks in the Business Service Workflow

The following sections provide information to help you manage the tasks that are created when requesting, changing, or deleting a business service:

- Section 16.1, “Displaying Administrator or Sponsor Tasks,” on page 187
- Section 16.2, “Claiming Tasks,” on page 187
- Section 16.3, “Approving and Denying Requests,” on page 188
- Section 16.4, “Completing Configuration Tasks,” on page 188
- Section 16.5, “Completing Trigger Tasks,” on page 190

For information about manually manipulating the task workflow, see Appendix C, “Using REST APIs to Customize Cloud Manager Behavior,” on page 205.

16.1 Displaying Administrator or Sponsor Tasks

Roles that Can Perform This Task: Cloud Administrator, Sponsor

In the *Unclaimed Tasks* list and the *All Tasks* list, tasks are categorized as Administrator tasks or Sponsor tasks. The Administrator tasks are the two configuration tasks (pre-build configuration and post-build configuration) and the Administrator approval task. The Sponsor task is the Sponsor approval task.

You can filter the two lists to show only Administrator tasks, only Sponsor tasks, or all tasks.

1. On the main navigation bar, click *Tasks*.
2. Click the list you want (*Unclaimed* or *All Tasks*).
3. In the select box at the top of the list, select the tasks you want to display (*All Tasks*, *Administrator Tasks*, or *Sponsor Tasks*).

16.2 Claiming Tasks

Roles that Can Perform This Task: Cloud Administrator, Approver (Administrator approval tasks only), Sponsor (Sponsor approval tasks only), Build Administrator (configuration tasks only)

You can claim an unclaimed task or a task that is owned by another user. When you claim a task, you become the task owner. No one else can work on the task unless you release (unclaim) the task or the other user claims it from you.

You cannot assign a task to another user. The task must be claimed by the other user.

1. On the main navigation bar, click *Tasks*.
2. To claim an unclaimed task, click the *Unclaimed* tab.

   or

   To claim a task owned by another user, click the *All Tasks* tab.
3. Select the task to claim, then click *Claim*.

   The task is added to your *My Tasks* list. It also remains in the *All Tasks* list but you are added as the task owner.

16.3 Approving and Denying Requests

Roles that Can Perform This Task: Cloud Administrator, Approver (Administrator approval tasks only), Sponsor (Sponsor approval tasks only)

As a Cloud Administrator, you can complete both Administrator approval tasks and Sponsor approval tasks.

1. On the main navigation bar, click *Tasks*, then click the *Unclaimed* tab.
2. Select the task for the business service request, then click one of the following options:
   - *Review*: Lets you review the business service before approving or denying the request.
   - *Approve*: Approves the request.
   - *Deny*: Denies the request. You must specify the reason for the denial. The user receives an e-mail with the reason and can modify the business service and resubmit the request.
   - *Claim*: Assigns the task to you and moves it to your *My Tasks* list. No other user can access the task. This is useful if you want to make sure that you own the task but want to complete it at a later time.
3. If you clicked *Review*, review or change the business service details, including the workload details. When you are finished, select one of the following:
   - *Approve*: Approves the request.
   - *Deny*: Denies the request. You must specify the reason for the denial. The user receives an e-mail with the reason and can modify the business service and resubmit the request.
   - *Claim*: Assigns the task to you and moves it to your *My Tasks* list.
   - *Close*: Closes the task without approving or denying it.

16.4 Completing Configuration Tasks

As a Cloud Administrator, you can complete both pre-build configuration and post-build configuration tasks.

- Section 16.4.1, “Completing Pre-Build Configuration Tasks,” on page 189
- Section 16.4.2, “Completing Post-Build Configuration Tasks,” on page 189

16.4.1 Completing Pre-Build Configuration Tasks

Roles that Can Perform This Task: Cloud Administrator, Build Administrator

The pre-build configuration task pauses the workflow process to give you an opportunity to perform any configuration required before the business service workloads are built. This might be configuration that you complete in the Cloud Manager console, such as configuring the Windows settings (Administrator account password, computer name, product license key, and registered name) or remote console password for one of the workloads. Or, it might be a task that you need to perform at the Cloud Manager Orchestration Server or in your hypervisor tools.

1. On the main navigation bar, click *Tasks*, then click the *Unclaimed* tab.
2. Select the task (the subject is *Pre-build configuration of a new Business Service* or *Pre-build configuration of a changed Business Service*).

   If there is configuration that must be completed for the business service in the Cloud Manager console, the *Complete* action is not available and an asterisk (*) appears next to it. You can mouse over the asterisk to view the remaining configuration requirements.
3. If you want to claim the task so that others cannot work on it, click *Claim*.

   Claimed tasks are moved to your *My Tasks* list.
4. Complete any configuration required for the workloads:
   - **Cloud Manager console tasks:** Select the configuration task, then click the *Review* action to display the task. Select a workload that needs to be configured, then click *Edit*. Configure the required settings and save the workload.
   - **External tasks:** Complete the tasks.
5. Mark the task as complete:
   - If you claimed the task, it is your *My Tasks* list. Select the task, then click *Complete*.
   - If you did not claim the task, it is in the *Unclaimed* list. Select the task, then click *Complete*.

16.4.2 Completing Post-Build Configuration Tasks

Roles that Can Perform This Task: Cloud Administrator, Build Administrator

The post-build configuration task pauses the provisioning workflow to give you an opportunity to perform any post-build configuration for the business service workloads. You must complete the configuration (if any) and then mark the task as completed to continue the workflow process.

1. On the main navigation bar, click *Tasks*, then click the *Unclaimed* tab.
2. If you want to claim the task so that no others can work on it, select the task, then click *Claim*.

   Claimed tasks are moved to your *My Tasks* list.
3. If you want to review the business service request, select the task (the subject is *Post-build configuration of a new Business Service* or *Post-build configuration of a changed Business Service*), then click *Review*.
4. Complete any configuration required for the workloads.

5 Mark the task as complete:

- If you claimed the task, it is your *My Tasks* list. Select the task, then click *Complete*.
- If you did not claim the task, it is in the *Unclaimed* list. Select the task, then click *Complete*.

16.5 Completing Trigger Tasks

As a Cloud Administrator, you can complete or schedule a reboot trigger task.

- Section 16.5.1, “Completing Reboot Trigger Tasks,” on page 190

16.5.1 Completing Reboot Trigger Tasks

Roles that Can Perform This Task: Business Service Owner, Cloud Administrator, Build Administrator

As a Cloud Administrator, you can complete any workload reboot trigger task. The trigger task pauses the provisioning change request workflow between the pre-build configuration task and before the changes are applied tot he workloads in the system build. This gives you an opportunity to immediately reboot or to schedule the reboot of a business service workload at a later time.

The reboot trigger task stops progress of the workflow, pending user interaction. The task blocks an unintentional reboot of a workload from occurring at an inopportune time. The *Reboot Now* feature in the task lets you decide when the system is in a state to accommodate a reboot, and then to take the appropriate action to reboot immediately.

The reboot scheduling feature accommodates situations when you want the workload to change significantly and these changes will result in the service being taken offline for a significant amount of time. Scheduling an off-hours, weekend, or vacation reboot can make the downtime less impactful to your business.

1 On the main navigation bar, click  *Tasks*, then click the *Unclaimed* tab.

2 Select the task for the business service change request (it should be tagged with a “needs reboot” message), then click one of the following options:

- *Review*: Lets you review the business service change request before immediately rebooting or scheduling a reboot of the workload.
- *Reboot Now*: Reboots the workload when the option is selected.
- *Schedule Reboot*: Lets you specify the date and time when you want the workload reboot to occur.
- *Claim*: Assigns the task to you and moves it to your *My Tasks* list. No other user can access the task. This is useful if you want to make sure that you own the task but want to complete it at a later time.
- *Unclaim*: Releases the task and returns it to the *Unclaimed* list.

3 If you clicked *Review*, you can review or change the business service details, including the workload details using one of these options:

- *Add*: Opens a list of workload templates where you can select an additional template to add and configure for the business service request.
- *Edit*: Lets you edit the configuration details of the selected workload.

- *Remove*: Deletes the selected workload from the business service change request.
- *Copy*: Lets you create copies of the selected workload.

You can also select *Claim*, *Reboot Now*, or *Schedule Reboot* from this Review Task dialog box, as explained in Step 2.

17 Generating Cloud Manager Reports

The following sections provide information to help you generate reports:

- Section 17.1, “Report Descriptions,” on page 193
- Section 17.2, “Generating Reports,” on page 194

17.1 Report Descriptions

Cloud Manager provides nine reports that provide information about business service costs, resource usage, organizations, and zones.

As a Cloud Administrator, you can generate all nine reports with an unlimited scope (all organizations, zones, business services, and so forth). Other roles can generate specific reports, as indicated in the following table, with a limited scope as determined by their roles.

For example, you can generate a Business Service Cost Details report for any organization, while an Organization Manager can generate the same report but only for his or her organization.

Report Name	Description
Business Service Cost Details	The cost details for each business service in an organization, organized by business group. Can be generated by: Cloud Administrator, Organization Manager, Sponsor, Business Service Owner, Business Group Viewer
Business Service Cost History	The history of all cost changes for each business service in an organization, organized by business group. Can be generated by: Cloud Administrator, Organization Manager, Sponsor, Business Service Owner, Business Group Viewer
Business Service Cost Summary	A summary of the setup, monthly, annual, and contract costs for all business services in an organization, organized by business group. Can be generated by: Cloud Administrator, Organization Manager, Sponsor, Business Service Owner, Business Group Viewer
Cloud Business Service Cost Details	The cost details for each business service in every organization, organized by organization and then business group. Can be generated by: Cloud Administrator
Organization Overview	A summary of the number of users, business groups, resource groups, business services, and workloads in the organization, along with the monthly and annual business service costs. Can be generated by: Cloud Administrator
Resource Group Details	A summary of the resource groups in a zone. Can be generated by: Cloud Administrator, Zone Administrator

Report Name	Description
Shared Storage Details	The details (such as capacity, used and free space, and stored VMs and templates) for all shared storage devices in a zone. Can be generated by: Cloud Administrator, Zone Administrator
Shared Storage Summary	A summary of the shared storage devices in a zone. Can be generated by: Cloud Administrator, Zone Administrator
Zone Overview	A summary of the number of resource groups, hosts, clusters, workloads, workload templates, networks, and storage devices in a zone. Can be generated by: Cloud Administrator, Zone Administrator

17.2 Generating Reports

Roles that Can Perform This Task: Cloud Administrator, Zone Administrator, Organization Manager, Sponsor, Business Service Owner, Business Group Viewer

You can generate reports in PDF, CSV, and XLS format. Generated reports are saved in your *My Reports* list until you delete them. For descriptions of the reports, see Chapter 17.1, “Report Descriptions,” on page 193.

To generate a report:

1. On the main navigation bar, click  *Reports*.
2. Click *Generate* to display the Reports dialog box.
3. In the *Report Templates* list, select the report you want to generate and the format you want, then click *Next*.
4. In the Report Parameters dialog box, select the organization or zone for which to generate the report, then click *Generate*.

   A report window appears. Depending on the amount of data to be collected, the report might be completed quickly or it might take a while. As soon as the report is completed, it is displayed in the report window, saved to your computer, opened in an associated application, or you are prompted about which action you want to take (depending on your browser configuration).

   If the report is taking a while, you can close the report window and the report continues to generate. If you close the report, its status is shown in the *My Reports* list. As soon as it is complete, you can view it.

18 Changing Application Server Default Parameters and Values

This section includes information you can use if you want to change the default parameters and values on your NetIQ Cloud Manager Application Server.

- Section 18.1, “Increasing the Java Heap Space on the Application Server,” on page 195

18.1 Increasing the Java Heap Space on the Application Server

Follow these steps if you want to increase the Java heap space on the NetIQ Cloud Manager Application Server.

1. At the Application server, open `/opt/netiq/cloudmanager/bin/setenv`,
2. Edit the file to change the heap space value:
   - **2a** Find the following line:

     ```
     #export JAVA_MAX_MEM # Maximum memory for the JVM
     ```

   - **2b** Change the line to look like this:

     ```
     export JAVA_MAX_MEM=4092M # Maximum memory for the JVM
     ```

3. Save and close the file.
4. Restart the Application server with this command:

   ```
   rcnetiq-cloudmanager restart
   ```

A Setting Up Cloud Manager to Log to a Sentinel Collector

NetIQ has created a Sentinel Collector to provide data capture capabilities for NetIQ Cloud Manager Application Server. Sentinel must be installed and operational before attempting to use this Collector.

The Collector parses, normalizes, and enhances records received from a data source (known as an Observer). Other Event Source Management (ESM) components like Connectors and Collector Managers perform functions such as remote protocol connections and data mapping. To learn more about Sentinel and its components, see the NetIQ Sentinel product page (http://www.netiq.com/products/sentinel/index.asp) or the NetIQ Sentinel product documentation (https://www.netiq.com/documentation/sentinel70/).

You can download a custom-built Sentinel Collector plug-in for Cloud Manager at the Sentinel Plug-ins Web site (http://support.novell.com/products/sentinel/secure/sentinelplugins.html). The site also has a link to download documentation for the Cloud Manager Collector.

If you choose to use this Collector, you need to configure Cloud Manager to send its syslog information to the Collector. Use the following steps to set up Cloud Manager.

1. At the Cloud Manager Application Server, modify the file: `/opt/netiq/cloudmanager/etc/cmauditlogger.properties`.
   - **1a** In the properties file, change the following line

     ```
     log4j.appender.CMSYSLOG.layout.ConversionPattern=%m\n
     ```

     to look like this:

     ```
     log4j.appender.CMSYSLOG.layout.ConversionPattern=NQ_CloudManager: %m\n
     ```
   - **1b** In the properties file, change the current audit location line

     ```
     log4j.category.com.novell.cm.audit.api.impl.AuditLogger=INFO, CMFILE
     ```

     to look like this:

     ```
     log4j.category.com.novell.cm.audit.api.impl.AuditLogger=INFO, CMFILE, CMSYSLOG
     ```
   - **1c** (Optional) If you don't want the local audit file, change the current audit location line

     ```
     log4j.category.com.novell.cm.audit.api.impl.AuditLogger=INFO, CMFILE
     ```

     to look like this:

     ```
     og4j.category.com.novell.cm.audit.api.impl.AuditLogger=INFO, CMSYSLOG
     ```
   - **1d** Save the properties file and restart the Application Server.
2. At the Cloud Manager Application Server, configure syslog to receive messages from the Application Server and then send it to the Sentinel server. To do this, modify the file: `/etc/syslog-ng/syslog-ng.conf`.
   - **2a** In the syslog file, add a new source. For example:

     ```
     source r_src { udp(ip("localhost") port(514)); };
     ```

2b (Conditional) If other services are already logging locally over UDP, you can add a filter line in the syslog file. For example:

```
filter f_ncm        { facility(syslog) and match('NQ_CloudManager:'); };
```

NOTE: The `syslog` value shown for the facility in the line above should match the value for the facility specified in the `cmauditlogger.properties` file. The default is `syslog`.

2c In the syslog file, create a destination and log entry for syslog. For example:

```
destination sentinel { tcp("###.###.###.##" port(1468)); };
log { source(r_src); filter(f_ncm); destination(sentinel); };
```

NOTE: The port number shown in the first line above must match the port for the Syslog TCP listener on the Sentinel server.

2d Save the file and restart syslog on the Application Server. On SUSE Linux Enterprise Server 11, the syslog restart command looks like this:

```
/etc/init.d/syslog-ng restart
```

3 Ensure that the Sentinel Collector for NetIQ Cloud Manager Collector is added to Sentinel and that the Syslog TCP connector in Sentinel is configured and running.

B Integrating Cloud Manager with the Nixu IP Address Management System

NetIQ Cloud Manager supports the integration of the Nixu NameSurfer Suite (version 7.2.3), an IP Address Management (IPAM) system that can simplify planning, monitoring, and managing the IP addresses in a network, including the addresses of various network interface cards on provisioned workloads.

Cloud Manager deploys its integration as an OSGI bundle with an accompanying configuration file. Use the steps in the following sections to enable Cloud Manager integration with an installed Nixu NameSurfer Suite.

- Section B.1, “Configuring Nixu NameSurfer for Cloud Manager Integration,” on page 199
- Section B.2, “Configuring Cloud Manager for Nixu NameSurfer Integration,” on page 201

B.1 Configuring Nixu NameSurfer for Cloud Manager Integration

In the Nixu NameSurfer environment, you need to create values to match specific Cloud Manager networking settings. This enables IPAM option at the network level.

The following steps assume you are familiar with the Nixu NameSurfer environment:

1. Create an NSAPI key to enable Nixu to communicate with the Cloud Manager Web Service.
   - **1a** From the NameSurfer *Configuration* menu, click *Keys* > *Add NSAPI key* to open the New NSAPI key view.
   - **1b** Specify a *Key name*, then save the *Key secret part* value for use later when you configure Cloud Manager for the IPAM integration.
   - **1c** (Optional) You can check the *Restrict access rights with groups* check box to restrict access rights associated with a defined Nixu user group.
2. Create a forward lookup zone to contain your IP address block.
   - **2a** From the NameSurfer *DNS* menu, click *Forward zones* > *Create zone*.
   - **2b** Configure the zone as required.
3. Add (or identify an existing) an IP address block to associate with each Cloud Manager network you will be using.
   - **3a** (Conditional) To add an IP address block, from the NameSurfer *IP Address* menu, click *Add block*, then configure the new block as required.

     or
   - **3b** (Conditional) From the NameSurfer *IP Address* menu, click *Root Blocks* to open the list of blocks, then select an unused block.

4 Configure the block with required DNS information for Cloud Manager integration.

4a From the *Blocks* list, select the block you want to configure, then select the *DNS* tab for this block.

4b Specify the *DNS zone information*. This is the name you gave to the forward lookup zone in Step 2.

4c (Optional) If you want Nixu to automatically generate a host name for the block, follow the pattern listed in the interface and enter the Host name pattern.

NOTE: This value is co-dependent with a value in the Cloud Manager configuration file. See autoHostname below, for more information.

---

5 From the block details page, select the *Information* tab, then select the *Add information* menu option to open an Information dialog box.

6 Create the following name/value pair in the Information dialog box for this block:

- **Name:** `NCMNetworkID`

  This is the Cloud Manager network ID, as displayed in the upper left-hand corner of the Edit Network dialog of the Cloud Manager Application Server Web console. Note that Cloud Manager configures this name without a period (.) in the string.

  Value: (example) `362:Network:digitalAirlines-Prod`

  This value is co-dependent with a value in the Cloud Manager configuration file. If you have set value of the `NCMNetworkIDValueIsName` property to TRUE in the configuration file, you need to provide only the Cloud Manager network name (for example, `DigitalAirlines-Prod`) instead of the Cloud Manager network ID.

7 Continue configuring the block to return the values that, if matched, are used to configure a corresponding Cloud Manager vNIC.

Open and complete information for each of the following name/value pairs:

- **Name:** `NCM.DNSServers`

  Value: (example) `192.168.0.1` 192.168.0.2

  What you specify for this value depends on your individual network setup. More than one value can be listed, with each value separated by a space.

- **Name:** `NCM.DNSSuffixes`

  Value: (example) `acme.com test.com da.com`

  What you specify for this value depends on your individual network setup. More than one value can be listed, with each value separated by a space.

- **Name:** `NCM.Gateways`

  Value: (example) `192.168.0.3` 192.168.0.4

  What you specify for this value depends on your individual network setup. More than one value can be listed, with each value separated by a space.

- **Name:** `NCM.NetMask`

  Value: (example) `255.355.255.0`

  What you specify for this value depends on your individual network setup.

B.2 Configuring Cloud Manager for Nixu NameSurfer Integration

1 From the Cloud Manager Application Server, open `/opt/netiq/cloudmanager/etc/cloudmanager-ipam-nixu.cfg`

2 Edit the configuration file:

2a Edit the properties that specify the Nixu Server connection information:

Configuration Property and Default Value	Value Detail	Example
`url=<server:port>/SOAP`	Specify the Nixu Server URL, without the "http" connection protocol.	`url=mynixuserver.mydomain.com:8443/SOAP/`
`secureConnection=true`	Specify whether the connection is secure. If http, specify `false`. If https, specify `true`.	`secureConnection=true`
`keyname=<keyname>`	This is the name of the NSAPI key configured in the Nixu NameSurfer console. It is needed for Web Service API access.	`keyname=EngKey`
`keyvalue=<keyvalue>`	This is the value of the NSAPI key specified in "keyname" property, as configured in the Nixu NameSurfer console. It is needed for Web Service API access.	`keyvalue=KrqcCHySczwIFfshSMfrPoCESY3ip8Nh90egtrmouTg=`

2b Edit the properties that specify the Cloud Manager options:

Configuration Property and Default Value	Value Detail
`autoHostname=TRUE`	Specifies whether Nixu NameSurfer should auto-create a hostname for the requested address. Hostname auto-generation follows the rules specified on the Address Block in the Nixu NameSurfer console.
`addToDNS=FALSE`	Specifies whether Nixu NameSurfer should auto-register the hostname with Nixu NameSurfer's DNS service. If you are not using Nixu NameSurfer for your DNS server, set this to `FALSE`.
`defaultToDHCP=FALSE`	Specifies behavior in the case of a Nixu NameSurfer integration failure. Failures can occur because of inability to connect to the server, inability to find a Nixu Address Block that is correctly configured to provide an address for a specified Cloud Manager Zone, etc. In case of a failure, if this value is `TRUE`, the Cloud Manager workload is configured to obtain its network information from DHCP.
`NCMNetworkIDValueIsName=FALSE`	Specifies whether to use the full Cloud Manager Network ID or the simple Network Name as the match string in Nixu NameSurfer Address Blocks. Using the full Cloud Manager Network ID is preferable, because it provides a greater probability of uniqueness. **NOTE:** This property is codependent with a Nixu IP Address block setting. See `NCMNetworkID`, above, for more information. Do not set this value to `TRUE` unless you are sure that all network interfaces have unique names.
`defaultVMNameFromHostName=FALSE`	Specifies whether to set the workload's `PreferredVMName` from the hostname provided by Nixu NameSurfer. We recommend that you set this value to FALSE, because the VM's name is now attached to a network name or IP address, which could change in the future if the Cloud Manager Network is deleted or modified. If TRUE, the `PreferredVMName` that will be set (the filename and display name of the VM) follows the pattern: <*hostname*>.<*workloadid*> **NOTE:** The `PreferredVMName` does not change in change requests, even if the hostname changes as a result of an altered NIC configuration.

Configuration Property and Default Value	Value Detail
`failBusinessServiceRequestOnIPAMFailure=TRUE`	Specifies what Cloud Manager should do in the case of a failure to obtain a valid NIC configuration from the IPAM system. the case of such a failure. If the value is set to TRUE, the business service is placed in a BUILD_FAILED state and the administrator can reconfigure the NIC and resubmit the business service. If the value is set to FALSE, the business service advances to the pre-build configuration phase. The administrator can then override the network type and specify a different NIC configuration.
`nixuIntegrationEnabled=FALSE`	Leave the value at FALSE until you finish setting all of the properties above. Setting the value to TRUE enables Cloud Manager to detect the Nixu IPAM integration and to display it as a network option.

3 Restart the Cloud Manager Application Server to refresh the configuration settings and enable IPAM integration.

C Using REST APIs to Customize Cloud Manager Behavior

This section includes the following information:

- Section C.1, “Cloud Manager REST API Overview,” on page 205
- Section C.2, “Using the REST APIs to Modify the Cloud Manager Workflow,” on page 205

C.1 Cloud Manager REST API Overview

NOTE: This section is not intended to explain REST concepts. It assumes that the reader has a prior understanding of REST concepts.

NetIQ Cloud Manager provides a complete set of REST APIs that are available for scripting or customizing the product. The APIs support both application/xml and application/json values for the Accept and Content-type request and response headers. You can use the GET, PUT, POST, and DELETE methods in the Cloud Manager REST calls as appropriate

You can visit `http://<server>:<port>/cloudmanager-api` for a full reference of the REST API, including the contents of the API packets include on return or should include on input.

C.2 Using the REST APIs to Modify the Cloud Manager Workflow

A Cloud Manager user can create, change, or delete a business service. Cloud Manager processes each of these three “requests” in its own unique workflow process: a series of steps occurring in the build and provisioning process.

During the workflow process, actions are taken sequentially as needed so that the process flows correctly. In general, the workflow proceeds in the order listed in the following table:

Table C-1 *Cloud Manager Workflow Phases*

Workflow Phase (Action Type)	Operation Type	Default, TaskBased Callout	Customizeable by Using Callouts
1. Submit	n/a	no	no
2. Approval	Approval	yes	yes
3. Pre-build configuration (also known as “PreConfig”)	Configuration	yes	yes
4. Mid approval	Approval	yes	yes
5. VM/workload build	n/a	no	no

Workflow Phase (Action Type)	Operation Type	Default, TaskBased Callout	Customizeable by Using Callouts
6. Post-build configuration (also known as “PostConfig”)	Configuration	yes	yes
7. Post-build approval (also known as “PostApproval”)	Approval	yes	yes
8. Pre-deploy (or “PreDeploy”)	Notification	no	yes

At the appropriate time in the workflow, the Cloud Manager Application Server invokes a uniquely registered and configured external process to execute each individual, categorized step. After the process (also known as a “callout”) performs its function in order, it reports its status to the Application Server, so that the server can invoke the next process. For Approval type operations, the completion status is either *Approved* or *Denied*. For Configuration type operations, the completion status is either *Success* or *Failure*. For Notify type operations, the callout is invoked asynchronously, so the server does not need to wait for confirmation. Without waiting for confirmation, the server signals this callout that it is time to execute its functions.

The default callout registered for each of these phases (except for PreDeploy) is handled by a task system (a TaskBasedCallout). One or more tasks are created for each phase and assigned, based on a set of rules. In theTaskBasedCallout system, each callout is configured with predefined tasks. For each business service request, the callout detects the tasks that must be created and to which Cloud Manager user (based on assigned roles and rights) the tasks should be assigned.

The workflow configuration can be specified at the system, organization, or business group layer. with the lowermost level dominating.

As a Cloud Administrator, you can skip any of the phases of the workflow and you can change or override default tasks by altering the permission assignment or the order.

As the Cloud Administrator, you can customize the workflow to invoke alternative callouts in place of Cloud Manager’s default task-based callouts. This might be useful if you want to integrate the Cloud Manager workflow with other existing approval processes or scripted configuration procedures. This section explains how to use REST calls to customize the callouts you need to integrate Cloud Manager’s workflow with the workflow of a 3rd party system.

This section includes the following information:

- Section C.2.1, “Workflow Callout Registration Concepts,” on page 206
- Section C.2.2, “Skipping the PreConfig Approval Stage by Invoking REST Callouts,” on page 207
- Section C.2.3, “REST URLs for Configuring and Registering Custom Callouts,” on page 208
- Section C.2.4, “Overriding the Default Task Order,” on page 212

C.2.1 Workflow Callout Registration Concepts

After you have created a custom callout, you can invoke a REST API call to actually register a callout in Cloud Manager. The registration defines a combination of conditions that, when satisfied, trigger Cloud Manager to invoke that callout. The conditions include the following:

- The domain level where the operation should happen, such as the entire Cloud Manager system, one of the organizations in the system, or one of the business groups in an organization.

When you register a callout, you do so at a domain level (system, organization, or business group). The registration levels are hierarchical. That is, if you do not register the callout at the organization level, it inherits the registration from its parent. You can register only one callout per operation, per domain level.

- The intended behavioral differences for the callout when applied to different business service request types:
  - Create a new business service (NEW)
  - Change an existing business service (CHANGE)
  - Delete an existing business service (DELETE)
- The intended action of “step” of the callout in the workflow, which would be one of the following:
  - Approval
  - PreConfig
  - PostConfig
  - PostApproval
  - PreDeploy

C.2.2 Skipping the PreConfig Approval Stage by Invoking REST Callouts

If you want to bypass (that is, “skip”) PreConfig approval of a workload that you customize for a user, it is not possible to do so using the Application Server Console, but you can invoke REST callouts that let you modify the skip flag in that stage of the workflow for new or changed business service requests.

The following skip flags are supported in Cloud Manager:

- skipNewWorkflowApproval
- skipNewWorkflowPreConfig
- skipNewWorkflowPostConfig
- skipNewWorkflowPostApproval
- skipChgReqWorkflowApproval
- skipChgReqWorkflowPreConfig
- skipChgReqWorkflowMidApproval
- skipChgReqWorkflowPostConfig
- skipChgReqWorkflowPostApproval
- skipDelWorkflowApproval
- skipDelWorkflowPreConfig

The following instructions for modifying PreConfig workflow REST callouts assume that you want to invoke the change for the entire Cloud Manager system. The instructions also assume that you are using the `application/json` payload type, although Cloud Manager also supports the `application/xml` payload type.

1. Invoke a REST GET call at the following URL:

   `http://<server>:<port>/8183/cloudmanager-api/system/integrationconfigs`

   The call returns the current settings. For example:

```
{
"skipChgReqWorkflowPostApproval": true,
"skipDelWorkflowPreConfig": true,
"skipNewWorkflowPostApproval": true,
"sourceObjId": 1,
"sourceObjName": "System",
"sourceObjType": "SYSTEM",
"url": "http://localhost:8183/cloudmanager-api/system/integrationconfigs"
}
```

The skip flags are tri-state: `true`, `false`, or `empty`. The empty state implies that it is inherited from the parent (if there is one). The `true` and `false` states imply an override of the parent setting, if there is one.

2 In the callout, add the skip flag for the PreConfig stage of the workflow when business services are created or changed:

```
{
"skipChgReqWorkflowPostApproval": true,
"skipDelWorkflowPreConfig": true,
"skipNewWorkflowPostApproval": true,
"skipNewWorkflowPreConfig": true,
"skipChgReqWorkflowPreConfig": true,
"sourceObjId": 1
}
```

In this example, the URL and the source object information are removed from the callout packet because they are not necessary for the PUT.

3 Invoke a REST PUT call at the same URL,

`http://<server>:<port>/8183/cloudmanager-api/system/integrationconfigs`

making sure that the payload type is correct, then save the PUT.

4 Confirm that the PUT was correctly by invoking a GET at the same URL and checking the packet structure, shown in Step 2.

NOTE: If you wanted to make this modification for a single system, you would replace the `http://<server>:<port>/8183/cloudmanager-api/system/integrationconfigs` URL with `http://<server>:<port>/8183/cloudmanager-api/organizations/<orgId>/integrationconfigs`.

C.2.3 REST URLs for Configuring and Registering Custom Callouts

You can configure Cloud Manager by using different REST API calls on the Cloud Manager Application Server. The REST Client plugin for Firefox, or the Advanced Rest Client plugin for the Google Chrome browser have been tested for such calls.

You can visit `http://<server>:<port>/cloudmanager-api` for a full reference of the REST API, including the contents of the API packets include on return or should include on input.

The tables below show the REST endpoints you can use to view the current Cloud Manager callouts. The URLs in the tables contain variables (that is, address elements surrounded by brackets < >) that imply specific information that you must provide:

- `/<domID>/` implies that you substitute one of the following:
  - `/system/`
  - `/organizations/<orgID>/`
  - `/businessgroups/<bgID>/`

Table C-2 *REST Calls for installed or available callouts*

URL Suffix	Method	Purpose of the Command / The Resulting Payload

`/callouts`	GET	**Purpose:** Represents the list of services mentioned in Section C.2, "Using the REST APIs to Modify the Cloud Manager Workflow," on page 205.
		
		Payload:
		
		`{"count":"4",`
		`"items":[`
		`{`
		`"bundleName":"com….taskBasedCallouts",`
		`"canonicalClassName":"org.eclipse….ServiceReferenceImpl",`
		`"implementedInterface":"interface com….ApprovalCalloutIfc",`
		`"provides":"Approval",`
		`"registrationName":"TaskBasedApproval",`
		`"simpleClassName":"ServiceReferenceImpl",`
		`"url":"http://host:8183/cloudmanager-api/callouts/TaskBasedApproval"`
		`},`
		`…`
		`] }`

`/callouts/{`*`name`*`}`	GET	**Purpose:** Gets the specified callout available in the OSGI system.
		
		{*name*}=`registrationName`
		
		NOTE: Everything in this payload is defined by the OSGI bundle, nothing is configurable, so no PUT operation is used.
		
		Payload:
		
		`{`
		`"bundleName":"com….taskBasedCallouts",`
		`"canonicalClassName":"org.eclipse….ServiceReferenceImpl",`
		`"implementedInterface":"interface com….ApprovalCalloutIfc",`
		`"provides":"Approval",`
		`"registrationName":"TaskBasedApproval",`
		`"simpleClassName":"ServiceReferenceImpl",`
		`"url":"http://host:8183/cloudmanager-api/callouts/TaskBasedApproval"`
		`}`

Table C-3 *REST Calls for integration configuration flags*

URL Suffix	Method	Purpose of the Command / The Resulting Payload

/<*domID*>/ integrationconfigs	GET	**Purpose:** Gets the domain object's integration configuration flags.
		Payload:
		{
		"skipChgReqWorkflowApproval":false,
		"skipChgReqWorkflowPostApproval":true,
		"skipChgReqWorkflowPostConfig":false,
		"skipChgReqWorkflowPreConfig":false,
		"skipDelWorkflowApproval":false,
		"skipDelWorkflowPreConfig":true,
		"skipNewWorkflowApproval":false,
		"skipNewWorkflowPostApproval":false,
		"skipNewWorkflowPostConfig":false,
		"skipNewWorkflowPreConfig":false,
		"sourceObjId":1,
		"sourceObjName":"System",
		"sourceObjType":"SYSTEM",
		"url":"http://host:8183/cloudmanager-api/system/integrationconfigs"
		}

/<domID>/ integrationconfigs	PUT	**Purpose:** Sets the domain object's integration configuration "skip" flags. **Payload:** Same payload as specified above. Skip flags have a tri-state: True, False, or empty. An empty value implies inheritance from the domain object's parent.
/<*domID*>/ integrationcalloutitems	GET	**Purpose:** Lists all callout registrations for the specified domain object. **Payload:** { "count":10, "items":[{ "key":"NEW_BUSINESS_SERVICE_REQUEST:POST_CONFIG", "sourceObjId":1, "sourceObjName":"System", "sourceObjType":"SYSTEM", "value":"TaskBasedPostConfig", "url":http://host:8183/cloudmanager-api/system/integrationcalloutitems/<*key*> }, …] }

URL Suffix	Method	Purpose of the Command / The Resulting Payload

/<*domID*>/integrationcalloutitems/{key}	GET	**Purpose:** Gets the specified callout registration.
		
		Payload: The payload in the example below indicates that at the system level, the `TaskBasedPostConfig` callout is used for PostConfig phase of new business service requests.
		
		{
		"key":"NEW_BUSINESS_SERVICE_REQUEST:POST_CONFIG",
		"sourceObjId":1,
		"sourceObjName":"System",
		"sourceObjType":"SYSTEM",
		"value":"TaskBasedPostConfig",
		"url":http://host:8183/cloudmanager-api/system/integrationcalloutitems/<*key*>
		}

/<*domID*>/integrationcalloutitems	PUT	**Purpose:** Creates or modifies the registration for the specified domain object.
		
		Payload:
		
		{
		"key":"NEW_BUSINESS_SERVICE_REQUEST:POST_CONFIG",
		"value":"TaskBasedPostConfig",
		}

/<*domID*>/integrationconfigitems	GET	**Purpose:** Gets the callout-specific configurations for the specified domain object. The “value” tag of the payload is specific to the callout.
		
		Payload: The payload in this example indicates that when the `TaskBasedPostConfig` callout is invoked for new business service requests. The callout creates a single task and assigns it to the TASK_COMPLETE_CONFIG_BS role.
		
		{
		"count":13,
		"items":[
		{
		"key":"TaskBasedPostConfig:NEW_BUSINESS_SERVICE_REQUEST",
		"sourceObjId":1,
		"sourceObjName":"System",
		"sourceObjType":"SYSTEM",
		"value":"1=TASK_COMPLETE_CONFIG_BS",
		"url":http://host:8183/cloudmanager-api/system/integrationconfigitems/key
		}
		…
		] }

URL Suffix	Method	Purpose of the Command / The Resulting Payload

/<*domID*>/integrationconfigitems/{key}	GET	**Purpose:** Gets the callout-specific configuration for the specified key and domain object. The “value” tag of the payload is specific to the callout.
		Payload:
		{
		"key":"TaskBasedPostConfig:NEW_BUSINESS_SERVICE_REQUEST",
		"sourceObjId":1,
		"sourceObjName":"System",
		"sourceObjType":"SYSTEM",
		"value":"1=TASK_COMPLETE_CONFIG_BS",
		"url":http://host:8183/cloudmanager-api/system/integrationconfigitems/key
		}

/<*domID*>/integrationconfigitems/{key}	PUT	**Purpose:** Sets or modifies the callout-specific configuration for the specified key and domain object. The “value” tag of the payload is specific to the callout.
		Payload:
		{
		"key":"TaskBasedPostConfig:NEW_BUSINESS_SERVICE_REQUEST",
		"value":"1=TASK_COMPLETE_CONFIG_BS",
		"url":http://host:8183/cloudmanager-api/system/integrationconfigitems/key
		}

C.2.4 Overriding the Default Task Order

The following is an example of how to use the REST API to get the current configuration:

method: GET

URL: `http://localhost:8183/cloudmanager-api/system/integrationconfigitems/TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST`

Return packet:

```
{
"key": "TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST",
"sourceObjId": 1,
"sourceObjName": "System",
"sourceObjType": "SYSTEM",
"value": "2=TASK_SPONSOR_BS@!@1=TASK_ADMIN_APPROVE_BS",
"url": "http://localhost:8183/cloudmanager-api/system/integrationconfigitems/TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST"
}
```

This GET call returns the current configuration for the specified configuration item. In this case, the specified item called for is the `TaskBasedApproval` step in new business services requested throughout the Cloud Manager system.

NOTE: You could replace the `/system/` component of the URL with `/organizations/<orgID>/` if you wanted to work with the configuration for a specific organization.

If a configuration doesn't exist at `/organizations/<orgID>/`, it inherits that configuration from the system level.

To change the current configuration, execute a PUT call with an appropriate packet. In the example to follow, assume a new custom permission, `MY_CUSTOM_PERMISSION`, has been created in the system and that specific users have been given this permission.

If you want to change the approval task order so that it is routed first to administrator approval permission and then to the new custom permission, you would use the following procedure:

Method: PUT

URL: `http://localhost:8183/cloudmanager-api/system/integrationconfigitems/TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST`

Input Packet:

```
{
"key": "TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST",
"sourceObjId": 1,
"value": "2=MY_CUSTOM_PERMISSION@!@1=TASK_ADMIN_APPROVE_BS"
}
```

In both of these sample REST calls, the PUT and the GET methods both use the `cloudmanager-api` configuration URL. The `value` property in the packet is defined by the specific callout. Task callouts expect an order of permissions, but custom callouts define their own value format.

To accomplish the same GET and PUT that we did above, but in a Task callout-specific way, use the following procedure, with the GET call first:

Method: GET

URL: `http://localhost:8183/tasks-api/system/taskconfigurations/TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST`

Return Packet:

```
{
  "calloutName": "TaskBasedApproval",
  "configItem": {
    "key": "TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST",
    "sourceObjId": 1,
    "sourceObjName": "System",
    "sourceObjType": "SYSTEM",
    "value": "2=TASK_SPONSOR_BS@!@1=TASK_ADMIN_APPROVE_BS"
  },
  "key": "TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST",
  "opType": "NEW_BUSINESS_SERVICE_REQUEST",
  "option": "",
  "taskAssignmentOrder":
  [
    {
      "name": "TASK_ADMIN_APPROVE_BS"
    },
    {
      "name": "TASK_SPONSOR_BS"
    }
  ],
  "url": "http://localhost:8183/tasks-api/system/taskconfigurations/TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST"
}
```

The GET call results in more information, and makes sense of the `value` property in a way that the Task Callout understands it.

The PUT call would look like this:

Method: PUT

URL: `http://localhost:8183/tasks-api/system/taskconfigurations/TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST`

Input Packet:

```
{
  "key": "TaskBasedApproval:NEW_BUSINESS_SERVICE_REQUEST",
  "taskAssignmentOrder":
  [
    {
      "name": "TASK_ADMIN_APPROVE_BS"
    },
    {
      "name": "MY_CUSTOM_PERMISSION"
    }
  ]
}
```

D Documentation Updates

This section contains information on documentation content changes that were made in this *Procedures Guide* after the initial release of NetIQ Cloud Manager 2.3. The changes are listed according to the date they were published.

NOTE: This updated information does not appear in the help content accessible from the product's user interface.

The documentation for this product is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the changes listed in this section.

If you need to know whether a copy of the PDF documentation that you are using is the most recent, the PDF document includes a publication date on the title page.

The documentation was updated on the following dates:

- Section D.1, "December 15, 2014," on page 215
- Section D.2, "December 11, 2014," on page 215
- Section D.3, "October 15, 2014," on page 216

D.1 December 15, 2014

Updates were made to the following sections:

Location	Update
Section 11.2, "Enabling Automatic User Creation," on page 105	Retitled the section (and similarly-worded content inside this section) from *Providing Self Registration for Users* to more intuitively describe the process of automatically bringing users into the Cloud Manager Application Console.
Section 11.3, "Creating User Groups," on page 107	Retitled the subsections to distinguish automatic user creation from manual user creation.

D.2 December 11, 2014

Updates were made to the following sections:

Location	Update
Various locations.	Removed references to all hypervisor support except VMware vSphere.

D.3 October 15, 2014

Updates were made to the following sections:

Location	Update
“Using the Right-Click Menu for Provisioning Actions” on page 57.	Added content to clarify the difference between the Suspend action and the Pause action (at the Orchestration Server console).
“Starting, Suspending, or Shutting Down a Workload” on page 176.	Added content to clarify the Suspend action (at the Application Server console) for the deployed workload.