NetIQ® Cloud Manager

Installation and Upgrade Guide

July 2016

Legal Notice

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see https://www.netiq.com/company/legal/.

Copyright © 2016 NetIQ Corporation. All Rights Reserved.

Contents

About This Book and the Library 9

Part I Preparing for Cloud Manager Installation 11

1 Installation Checklist 13

2 Cloud Manager System Requirements 17

2.1 Cloud Manager Licensing Requirements . . . 17
2.1.1 License Requirements During Installation . . . 17
2.1.2 Product License . . . 17
2.1.3 Product Evaluation Terms and Procurement . . . 18
2.1.4 Upgrading the Server from a Trial License to a Purchased License . . . 18
2.2 Cloud Manager Orchestration Server Requirements . . . 18
2.2.1 Requirements For the Orchestration Server . . . 19
2.2.2 Port Requirements for the Orchestration Server . . . 19
2.3 Cloud Manager Orchestration Agent Requirements . . . 21
2.3.1 Requirements for Orchestration Agent . . . 21
2.3.2 Port Requirements for the Orchestration Agent . . . 22
2.4 Cloud Manager Orchestration Console Requirements . . . 22
2.5 Cloud Manager Application Server Requirements . . . 23
2.5.1 Requirements for the Application Server . . . 23
2.5.2 Port Requirements for the Application Server . . . 24
2.6 Cloud Manager Application Console Requirements . . . 25
2.7 Cloud Manager Monitoring Agent Requirements . . . 25
2.8 Requirements and Cloud Manager Support for the Virtual Environment . . . 26
2.8.1 RHEL 6 VM Support . . . 28
2.8.2 Requirements for Machines Designated as VM Hosts . . . 28

3 Choosing the Installation Patterns and Where to Install Them 29

3.1 Cloud Manager Orchestration Server Installation Pattern . . . 29
3.2 Cloud Manager Installation Pattern . . . 30
3.3 Cloud Manager Monitoring Server Installation Pattern . . . 31
3.4 Cloud Manager Orchestration Agent Installation Pattern . . . 31
3.5 Cloud Manager Monitoring Agent Installation Pattern . . . 32
3.6 Cloud Manager Orchestration Console Installation Pattern . . . 32

4 Preinstallation Tasks 35

4.1 Gathering Certificate and License Information . . . 35
4.2 Importing the Public Key for Cloud Manager . . . 35
4.3 Preparing the Server When Multiple NICs and DNS Addresses Exist . . . 36

Part II Installing Cloud Manager 37

5 Installing Cloud Manager Orchestration Components 39

5.1 Installing Orchestration Components on SLES . . . 40
5.2 Alternative Installation Methods for the Orchestration Agent . . . 41
5.2.1 Obtaining the Agent Installer and Supporting Files from the Administrator Information Page . . . 42
5.2.2 Installing the Agent on Windows Machines . . . 43
5.2.3 Manually Installing the Agent Packages on SLES Machines . . . 44
5.2.4 Manually Installing the Agent Linux Packages on RHEL Machines . . . 44
5.2.5 Advanced Agent Installation Methods . . . 45
5.3 Alternative Installation Methods for the Orchestration Console and Clients . . . 48
5.3.1 Obtaining Installers from the Administrator Information Page . . . 49
5.3.2 Installing the Console and Clients on Windows . . . 49
5.3.3 Installing the Console and Clients on a SLES Server . . . 51
5.4 Alternative Installation Methods for the Cloud Manager Monitoring Agent . . . 51
5.4.1 Installing the Cloud Manager Monitoring Agent on Linux Servers . . . 52
5.4.2 Installing the Cloud Manager Monitoring Agent On Windows Machines . . . 52

6 Installing Cloud Manager Application Server Components 55

6.1 Installing Cloud Manager Pattern on Your SLES Server . . . 55

Part III Configuring Cloud Manager 57

7 Configuring Cloud Manager Orchestration Components 59

7.1 Configuring the Orchestration Server . . . 60
7.2 Configuring the Monitoring Server and Monitoring Agent . . . 62
7.3 Configuring the Orchestration Agent . . . 62
7.4 Validating and Optimizing the Orchestration Installation . . . 64
7.5 Orchestration Server Connection and Read Timeout Values . . . 64

8 Configuring Connections to the Cloud Manager Application Server 67

8.1 Enabling a Secure Connection . . . 67
8.1.1 Configuring the Cloud Manager Web Service Secure Port . . . 67
8.1.2 Disabling SSL v3 in Java Jetty . . . 68
8.2 Enabling a Non-Secure Connection . . . 69

9 Launching the Orchestration Console and Logging in to the Orchestration Server 71

9.1 Launching the Orchestration Console . . . 71
9.2 Logging In Explicitly to a Named Server . . . 72

10 Creating a Resource Account 73

10.1 Opening the Resources Monitor . . . 74
10.2 Automatically Registering a Resource . . . 74
10.3 Selecting a Resource for Manual Registration . . . 75
10.4 Manually Registering a Resource in the Orchestration Console . . . 76
10.4.1 Using the Orchestration Console to Create a Resource Account . . . 76
10.4.2 Installing an Orchestration Agent to Match the New Resource . . . 77

11 Configuring the vSphere Provisioning Adapter 81

11.1 Initializing the vSphere Provisioning Adapter on the Orchestration Server . . . 81
11.2 Managing the vSphere Provisioning Adapter . . . 83
11.2.1 VM Management Policies for the vSphere Provisioning Adapter . . . 83
11.2.2 Synchronizing vSphere vCenter System Changes to the Orchestration Server . . . 84
11.2.3 Assigning a vSphere VM to a Resource Pool . . . 84
11.2.4 Setting Up Orchestration VNC for a VM Managed by vSphere . . . 85
11.3 Configuring the vSphere Provisioning Adapter for VMware DRS Clusters and MetroClusters . . . 87
11.3.1 Setting Up Orchestration to Accommodate VMware DRS Clustering and Updates . . . 88
11.3.2 Constraining vSphere VMs to Their Assigned Resource Pools . . . 88
11.3.3 Configuring the Orchestration Server to Limit Datastore Visibility in vSphere Clusters . . . 88
11.3.4 Configuring Cloud Manager Orchestration for a vSphere MetroCluster Environment . . . 89
11.4 Discovering Enterprise Resources in Multiple vSphere Environments . . . 90
11.4.1 Creating Accounts for Each vCenter Environment . . . 90
11.4.2 Configuring the vsphere.vcenters Fact to Include All Accounts Representing a vCenter Server . . . 91
11.4.3 Optionally Specifying an Authentication Certificate for Each vCenter Server . . . 92
11.5 Discovering vSphere Hosts and VMs . . . 93
11.5.1 Running Discovery for a vSphere Environment . . . 93
11.5.2 Running Discovery for Multiple Resource Environments . . . 94

12 Configuring vSphereUpdate Client and Monitoring Jobs 95

12.1 Understanding vSphereUpdate Monitor Jobs . . . 95
12.2 Configuring the vSphereUpdate Client . . . 97
12.3 Troubleshooting vSphereUpdater Connections to vCenter . . . 99

13 Configuring Sysprep or Autoprep 101

13.1 Understanding and Configuring Sysprep . . . 101
13.1.1 How Sysprep Works . . . 101
13.1.2 Setting Sysprep Facts in the Orchestration Console . . . 102
13.1.3 Using the Sysprep deploy.cab Files . . . 110
13.1.4 Applying Sysprep Facts . . . 112
13.1.5 Example Sysprep Scenarios . . . 113
13.1.6 Known Sysprep Limitations . . . 113
13.2 Understanding and Configuring Autoprep . . . 114
13.2.1 How Autoprep Works . . . 114
13.2.2 Setting Autoprep Facts in the Orchestration Console . . . 115
13.2.3 Applying Autoprep Facts . . . 118
13.2.4 Example Autoprep Scenarios . . . 119
13.2.5 Known Autoprep Limitations . . . 119

14 Using the Cloud Manager Application Server Configuration Tool 121

14.1 Configuring the PostgreSQL Database Connection and Credentials . . . 122
14.2 Configuring Auto-Vacuum Settings for Tables in the PostgreSQL Database . . . 125
14.3 Configuring Cloud Manager to Use Authentication Sources . . . 126
14.3.1 Configuring Authentication to an LDAP Directory . . . 127
14.3.2 Configuring Authentication to NetIQ Access Manager . . . 129
14.4 Installing and Configuring Other Cloud Manager Feature Settings . . . 130
14.4.1 Installing the Cloud Manager Application Console . . . 130
14.4.2 Configuring the Cloud Manager Web Server (Jetty) . . . 130
14.4.3 Configuring the Cloud Manager Web Server to Use SSL . . . 131
14.4.4 Configuring Cloud Manager SMTP Mail Settings . . . 131
14.4.5 Configuring Cloud Manager System Shell Login Information . . . 132

Part IV Advanced Installation and Integration Topics 133

15 Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support 135

15.1 Overview . . . 135
15.2 Orchestration Server Failover Behaviors . . . 136
15.2.1 Use Case 1: Orchestration Server Failover . . . 136
15.2.2 Use Case 2: VM Builder Behavior at Orchestration Server Failover and Failback . . . 137
15.2.3 Use Case 3: Monitoring Behavior at Orchestration Server Failover and Failback . . . 137
15.3 Installing the Orchestration Server to a SLE HAE Cluster Environment . . . 137
15.3.1 Meeting the Prerequisites . . . 137
15.3.2 Installing the SLE High Availability Extension . . . 139
15.3.3 Configuring Nodes with Time Synchronization and Installing Pacemaker to Each Node . . . 139
15.3.4 Setting Up the OCFS2 File System for High Availability . . . 141
15.3.5 Installing the Orchestration Server on the First Cluster Node . . . 142
15.4 Configuring the Orchestration Server for High Availability . . . 143
15.4.1 Some Considerations When Configuring with the GUI Wizard . . . 144
15.4.2 The Configuration Procedure . . . 144
15.4.3 Checking the Configuration . . . 146
15.4.4 Running the High Availability Configuration Script . . . 146
15.5 Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster . . . 147
15.6 Creating the Server Cluster Resource Group . . . 147
15.7 Testing the Failover of the Orchestration Server in a High Availability Grid . . . 148
15.8 Installing and Configuring other Orchestration Components to the High Availability Grid . . . 148

16 Configuring the Orchestration Server to Use an Audit Database 149

16.1 Installing the PostgreSQL Package and Dependencies on an Independent Host . . . 149
16.1.1 Detail . . . 150
16.2 Configuring PostgreSQL to Accept Remote Database Connections . . . 151
16.3 Logging in Locally to the PostgreSQL Database . . . 152
16.4 Creating an Orchestration Server User for the PostgreSQL Database . . . 152
16.5 Configuring the Orchestration Server Audit Database on a Separate Host . . . 153
16.6 Installing and Configuring the Orchestration Server for Use with a Local PostgreSQL Audit Database . . . 154
16.6.1 Installing the PostgreSQL Package and Dependencies . . . 154
16.6.2 Configuring PostgreSQL to Accept Local Database Connections . . . 155
16.6.3 Logging in Locally to the PostgreSQL Database . . . 155
16.6.4 Installing and Configuring the Local Orchestration Server Audit Database . . . 156
16.7 Configuring the Audit Database after the Cloud Manager Orchestration Server Is Configured . . . 157
16.8 Configuring the Remote Audit Database after the Cloud Manager Orchestration Server Is Configured . . . 158
16.9 Modifying Audit Database Tables to Accommodate Long Names . . . 159
16.10 Understanding Grid ID Usage in the Audit Database . . . 159

17 Integrating the Orchestration Server with a Sentinel Collector 161

17.1 Integration Architecture . . . 161
17.2 System Requirements . . . 162
17.3 Importing and Deploying the Orchestration Server Sentinel Collector Plug-in . . . 163
17.4 Connecting the Orchestration Server to the Sentinel Collector Plug-In . . . 165
17.5 Verifying the Sentinel Configuration After Connecting to the Orchestration Server . . . 165
17.6 Event Classification and Taxonomy Keys . . . 166

17.7 Plain Text Visibility of Sensitive Information 169

18 Configuring Secure Authentication Sources to Communicate with Cloud Manager 171

18.1 Configuring NetIQ Access Manager to Work with Cloud Manager. 171
18.1.1 Managing a Reverse Proxy for Authentication to Cloud Manager 171

Part V Upgrading Cloud Manager 179

19 Upgrade Checklist 181

20 Upgrading the Cloud Manager Application Server Components 183

20.1 Backing Up the PostgreSQL Database 183
20.2 Backing Up the Application Server Components and Custom Files. 184
20.3 Install the Cloud Manager Package 185
20.4 Running the Cloud Manager Configuration Script 185
20.5 Using the Configurator Tool to Update Resource Pool Data on the Cloud Manager Application Server 186

21 Orchestration Components Upgrade Overview 187

21.1 Basic Functions of the Orchestration Components Upgrade 187
21.2 Cloud Manager Orchestration Components That Are Not Upgraded. 188

22 Upgrading Orchestration Server Components 189

22.1 Backing Up Orchestration Components Prior to Upgrading. 189
22.2 Checking the Current Version of Cloud Manager Orchestration Components. 190
22.3 Snapshotting the Existing Orchestration Server Installation 190
22.4 Upgrading Orchestration Components. 191
22.4.1 Upgrading Orchestration Packages Using YaST2 191
22.4.2 Upgrading Orchestration Packages Using the zypper Command 193
22.4.3 Other Useful zypper Commands for Upgrade 194
22.4.4 Checking the Upgraded Version of the Orchestration Components 194
22.5 Configuring the Upgraded Orchestration Components 195
22.5.1 Some Considerations When Configuring with the GUI Wizard 195
22.5.2 Configuring the Upgraded Orchestration Server 195
22.5.3 Configuring the Upgraded Monitoring Server 198
22.5.4 Configuring an Upgraded Orchestration Agent Installed on the Server. 198
22.6 Configuring the Remote Audit Database after Orchestration Components Are Upgraded. 199
22.7 Running Discovery on VM Hosts and Images 200
22.8 Alternate Methods for Upgrading Older Agents and Clients 200
22.8.1 Automatically Upgrading the Orchestration Agent from the Cloud Manager Orchestration Console 200
22.8.2 Using the ISO to Upgrade the Orchestration Agent on Red Hat Enterprise Linux 6 Machines 201
22.8.3 Using the ISO to Upgrade the Old Orchestration Agent or the Orchestration Clients on Windows Machines. 202
22.8.4 Using the Administrator Information Page to Upgrade the Agents and Clients 202
22.9 Running the Upgrade Configuration on an Enterprise Scale 202
22.10 Upgrading a Orchestration High Availability Configuration. 203

A Restoring Cloud Manager Application Server in the Event of a System Failure 205

B Compatibility Checking Behavior for Orchestration Components 207

B.1 If the Orchestration Server Is Not Compatible with the Orchestration Console 207
B.2 When an Agent Version Does Not Match the Server Version 208

C Recovering from a Failed Orchestration Server Upgrade 211

C.1 Upgrade Failure Scenarios 211
C.1.1 Failure Scenario 1: Error Resolution 211
C.1.2 Failure Scenario 2: Cannot Resolve Error 211

Part VI Uninstalling Cloud Manager 213

23 Uninstalling Orchestration Component Patterns from a SLES Server 215

D Documentation Updates 217

D.1 July 2016 217

About This Book and the Library

The *Installation and Upgrade Guide* provides detailed planning and installation information for the Cloud Manager components, as well as upgrade information.

- Part I, “Preparing for Cloud Manager Installation,” on page 11
- Part II, “Installing Cloud Manager,” on page 37
- Part III, “Configuring Cloud Manager,” on page 57
- Part IV, “Advanced Installation and Integration Topics,” on page 133
- Part V, “Upgrading Cloud Manager,” on page 179
- Part VI, “Uninstalling Cloud Manager,” on page 213

Intended Audience

This document is intended for IT staff, such as IT service providers and operators, who are responsible for deploying and configuring Cloud Manager.

Information in the Library

The library for this product is available in HTML and PDF formats on the Cloud Manager Documentation (https://www.netiq.com/documentation/cloud-manager) website.

The Cloud Manager library provides the following information resources:

Release Notes

Provides information about new features and enhancements in the release, as well as any known issues.

Product Overview

Provides information about the features, functionality, and operational concepts of Cloud Manager.

Installation and Upgrade Guide

Provides detailed planning and installation information for the software, as well as upgrade information.

Procedures Guide

Provides conceptual information, an overview of the user interface, and step-by-step guidance for common tasks.

Administrator Reference

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

Help

Provides context-sensitive information and step-by-step guidance for common tasks as you work in the interface.

Feedback

If you have suggestions for documentation improvements, click **comment on this topic** at the bottom of any page in the HTML version of the documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

Additional Resources

We encourage you to use the following additional resources online:

- Cloud Manager Forum (https://forums.netiq.com/forumdisplay.php?13-Cloud-Manager): A web-based community of product users where you can discuss product functionality and advice with other product users.
- Cloud Manager Product page (https://www.netiq.com/products/cloud-manager/): A web-based product brochure that provides information about features, how to buy, technical specifications, frequently asked questions, and a variety of resources such as videos and white papers.
- NetIQ User Community (https://www.netiq.com/communities/): A web-based community with a variety of discussion topics.
- NetIQ Support Knowledgebase (https://www.netiq.com/support/kb/): A collection of in-depth technical articles.
- NetIQ Support Forums (https://forums.netiq.com/forum.php): A web location where product users can discuss NetIQ product functionality and advice with other product users.
- MyNetIQ (https://www.netiq.com/f/mynetiq/): A website offering product information and services, such as access to premium white papers, webcast registrations, and product trial downloads.

Technical Support

For specific product issues, please contact our Technical Support team at www.netiq.com/support.

Preparing for Cloud Manager Installation

The information in this section can help you prepare your data center for the Cloud Manager installation.

- Chapter 1, “Installation Checklist,” on page 13
- Chapter 2, “Cloud Manager System Requirements,” on page 17
- Chapter 3, “Choosing the Installation Patterns and Where to Install Them,” on page 29
- Chapter 4, “Preinstallation Tasks,” on page 35

1 Installation Checklist

To ensure that you successfully install and configure Cloud Manager, you should follow the installation checklist provided below. Each task provides brief information and a reference to where you can find more complete details.

Status	Installation and Configuration Task	For information
❑	1. Review Cloud Manager concepts and terminology. Cloud Manager includes functionality and components you need to understand to successfully install, configure, maintain, and use the product. Cloud Manager also interacts with other products such as hypervisors and directory services.	If you are not already familiar with Cloud Manager concepts and its interaction with these other products, see the *Cloud Manager Product Overview*.
❑	2. Virtualize your physical datacenter. If you have not already applied the virtualization infrastructure to your physical datacenter, you need to do so.	See Section 2.8, “Requirements and Cloud Manager Support for the Virtual Environment,” on page 26.
❑	3. Review the supported Cloud Manager environments and software installation requirements.	See Chapter 2, “Cloud Manager System Requirements,” on page 17 and Section 2.8, “Requirements and Cloud Manager Support for the Virtual Environment,” on page 26.
❑	4. Prepare for Orchestration components installation. Some security certificate and licensing tasks must take place before you begin installing Orchestration components.	See Chapter 4, “Preinstallation Tasks,” on page 35.
❑	5. Install the Cloud Manager Orchestration Server, the Orchestration Console, and the Orchestration Agent. The Cloud Manager Orchestration Server communicates with its Orchestration Agents. These agents establish communication with your virtualization infrastructure (hypervisor technology). With this link in place, the Cloud Manager Orchestration Server utilizes specialized provisioning adapter jobs to automate the provisioning, management, and deprovisioning of virtual machines. There are some alternative methods you can use for installing these components.	See Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 39.

Status	Installation and Configuration Task	For information
❑	6. Configure the Cloud Manager Orchestration components. The Orchestration components and Application components are typically installed on different servers. You should configure the Orchestration components before you proceed with the Application installation: ◆ Configuring the Orchestration Server ◆ Configuring the Orchestration Agent ◆ Creating a resource account in the Orchestration Console ◆ Getting provisioning adapters running for VM discovery ◆ Configuring the Orchestration web service to connect to the Cloud Manager Application Server	See the following: ◆ Chapter 7, “Configuring Cloud Manager Orchestration Components,” on page 59 ◆ Chapter 10, “Creating a Resource Account,” on page 73 ◆ Chapter 11, “Configuring the vSphere Provisioning Adapter,” on page 81 ◆ Chapter 8, “Configuring Connections to the Cloud Manager Application Server,” on page 67 See the following for advanced functionality available in the Orchestration components: ◆ Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 41 ◆ Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 48
❑	7. Prepare for Cloud Manager installation. Before you install and configure Cloud Manager Application components, you need to prepare a remote database for storing Cloud Manager data. You must also prepare your authentication source to support Cloud Manager authentication. Cloud Manager supports the following authentication methods: ◆ NetIQ eDirectory (LDAP) ◆ Microsoft Active Directory (LDAP) ◆ NetIQ Access Manager	See Section 14.1, “Configuring the PostgreSQL Database Connection and Credentials,” on page 122. See Section 14.3, “Configuring Cloud Manager to Use Authentication Sources,” on page 126.
❑	8. Install the Cloud Manager Application Server and its console. The Cloud Manager Application Server provides the interface through which users request virtual resources. Requests are communicated to the Orchestration Server, which performs the required virtualization operations in conjunction with your hypervisor technology.	See Chapter 6, “Installing Cloud Manager Application Server Components,” on page 55.

Status	Installation and Configuration Task	For information
❑	9. Configure the Cloud Manager system. After installation, you must complete several configuration tasks before Cloud Manager can be used, including ◆ Configuring the Cloud Manager Application Server authentication source connections ◆ Configuring the Cloud Manager Application Server database connection	See Chapter 14, “Using the Cloud Manager Application Server Configuration Tool,” on page 121.

After you have completed the installation and configuration of both the Orchestration components and the Cloud Manager system, you are ready to start populating your Cloud Manager Application Server and Application Console with the components that enable users to provision their own business services. Continue with “Setting Up the Cloud Environment” in the *Cloud Manager Procedures Guide*.

2 Cloud Manager System Requirements

Before you install the Cloud Manager, ensure that your system resources meet the requirements of the product. This section includes the following information to help you with that evaluation:

- Section 2.1, “Cloud Manager Licensing Requirements,” on page 17
- Section 2.2, “Cloud Manager Orchestration Server Requirements,” on page 18
- Section 2.3, “Cloud Manager Orchestration Agent Requirements,” on page 21
- Section 2.4, “Cloud Manager Orchestration Console Requirements,” on page 22
- Section 2.5, “Cloud Manager Application Server Requirements,” on page 23
- Section 2.6, “Cloud Manager Application Console Requirements,” on page 25
- Section 2.7, “Cloud Manager Monitoring Agent Requirements,” on page 25
- Section 2.8, “Requirements and Cloud Manager Support for the Virtual Environment,” on page 26

2.1 Cloud Manager Licensing Requirements

Cloud Manager requires any installation of its Orchestration Server to be licensed. None of the Cloud Manager components functions properly without a licensed server installation.

- Section 2.1.1, “License Requirements During Installation,” on page 17
- Section 2.1.2, “Product License,” on page 17
- Section 2.1.3, “Product Evaluation Terms and Procurement,” on page 18
- Section 2.1.4, “Upgrading the Server from a Trial License to a Purchased License,” on page 18

2.1.1 License Requirements During Installation

The Cloud Manager Orchestration Server is normally the first component of the product to be installed. During the configuration of the downloaded server package, you are required to provide a path to a license file. This file can be either a license you purchase from NetIQ, or an evaluation license (a trial key). The installation configuration cannot proceed without the license.

2.1.2 Product License

To purchase a license for Cloud Manager, contact an authorized NetIQ Sales representative at 888-323-6768, or go to the Cloud Manager *How to Buy* web page (https://www.netiq.com/products/cloud-manager/how-to-buy/). Your representative will send a purchased license file for Cloud Manager to your account at the NetIQ Customer Center (https://www.netiq.com/customercenter). From the Customer Center, save the file to a location that you can access during the Cloud Manager Orchestration Server installation.

NOTE: If you do not have an eLogin account for the Customer Center, click here to create one.

For installation and configuration information, see Part II, “Installing Cloud Manager,” on page 37.

2.1.3 Product Evaluation Terms and Procurement

You can download and evaluate Cloud Manager with an evaluation license (trial key) for 90 days. The trial key controls the number of users and managed nodes you can configure, and sets an expiration date. After 90 days, you must purchase a license, or discontinue use of the product.

To initiate a 90-day trial for Cloud Manager, contact an authorized NetIQ Sales representative at 888-323-6768, or go to the Cloud Manager *How to Buy* web page (https://www.netiq.com/products/cloud-manager/how-to-buy/). Your representative will send a trial key file for Cloud Manager to your account at the NetIQ Customer Center (https://www.netiq.com/customercenter). From the Customer Center, save the file to a location that you can access during the Cloud Manager Orchestration Server installation.

NOTE: If you do not have an eLogin account for the Customer Center, click here to create one.

For installation and configuration information, see Part II, “Installing Cloud Manager,” on page 37.

2.1.4 Upgrading the Server from a Trial License to a Purchased License

If you are operating the Cloud Manager Orchestration Server with an evaluation license, use the following steps to upgrade from a trial key to a product license key you purchased from NetIQ:

1. Stop the Orchestration Server.

   ```
   /etc/init.d/netiq-cmosserver stop
   ```

2. Copy the purchased license key file (`key.txt`) to the `/opt/novell/zenworks/zos/server/license` directory. You overwrite an older license file in this process.

3. Start the Orchestration Server.

   ```
   /etc/init.d/netiq-cmosserver start
   ```

2.2 Cloud Manager Orchestration Server Requirements

Use the requirements in this section to plan your deployment of the Cloud Manager Orchestration Server.

- Section 2.2.1, “Requirements For the Orchestration Server,” on page 19
- Section 2.2.2, “Port Requirements for the Orchestration Server,” on page 19

2.2.1 Requirements For the Orchestration Server

The network machine where you install Cloud Manager Orchestration Server software must meet the requirements in Table 2-1:

Table 2-1 *Orchestration Server Requirements*

Item	Requirement
Operating System	SUSE Linux Enterprise Server 11 Service Pack 4 (SLES 11 SP 4) on the 64-bit (x86-64) architecture (Intel processor or AMD Opteron processor)
Hardware	◆ **Processor:** ◆ **Minimum:** 2.5 GHz 64-bit (or equivalent) Intel processor or AMD Opteron processor ◆ **Recommended:** Dual-Core, 2.5 GHz (or greater) 64-bit, Intel processor or AMD Opteron processor ◆ **RAM:** ◆ **Minimum:** 3 GB ◆ **Recommended:** 4 GB (or greater) ◆ **Disk Space:** ◆ **Minimum:** 350 MB for installing ◆ **Recommended:** 1 GB for managing fewer than 100 resources
Hostname Resolution	The server must resolve device host names by using a method such as DNS (recommended).
IP Address	The server must have a static IP address or a permanently leased DHCP address.

2.2.2 Port Requirements for the Orchestration Server

Table 2-2 identifies the ports and protocols required by the Cloud Manager Orchestration Server. Ports are configurable except where they are specified as mandatory. The Orchestration Server firewall must allow traffic for the ports (or their substitutes if you are not using the defaults). Network firewalls need to allow outgoing agent connections to TCP ports 8100 (not secure) and 8101 (secure) on the server.

Table 2-2 *Port Requirements for the Orchestration Server*

Port Number	Protocol	Function
80	TCP	(Not secure) Used for the Monitor Web Server. TCP Port 80 can be used by the Monitor Collector instead of TCP port 8649.
8001	TCP	Used for communication with the Administrator Information page.
8100	TCP	(Not secure) Used with a custom protocol for communication with the Orchestration Agent and for invoking the `zos` command line interface or opening the Java Developer's toolkit.

Port Number	Protocol	Function
8101	TCP	(Secure) Used for invoking the `zos` command line interface or opening the Java Developer's toolkit by using TLS.
8082	TCP	(Not secure) Used by the Cloud Manager Application Server to communicate with the Cloud Manager Orchestration Server. The Cloud Manager Orchestration Web Service exposes a RESTful interface used by the Cloud Manager Application Server to communicate with the Cloud Manager Orchestration Server through TCP ports 8082 and 8443.
8443	TCP	(Secure) Used by the Cloud Manager Application Server to communicate with the Cloud Manager Orchestration Server. See also port 8082.
8649	TCP and UDP	Used by monitored systems (physical and virtual) to send metrics to the Monitoring Server. The port must be available for gmond to communicate with the Orchestration Agent. The Monitoring Server is installed on the same system as the Cloud Manager Orchestration Server. TCP 8649 is used by the Monitor Collector component of the Monitoring Server. You can alternatively use TCP port 80. TCP 8649 is used on the monitored system to communicate to the localhost.
1099	TCP and UDP	**Mandatory**. Used with RMI for invoking the `zosadmin` command line interface or for running the Orchestration Console.
VNC ports (1000–1xxx)	TCP	Used by connections to VM consoles through a VNC client. Typically, this means TCP port 1000 for the first VM on a VM host, 1001 for the second, and so on. These connections go to the VM host, exposing the console on behalf of the VM.
4000	UDP	Used as a datagrid multicast request port and a control channel port.
4011–4014	UDP	Used for datagrid multicast file transfers. Enable for outbound communications. Multicast groups for datagrid multicast-based file transfers are 239.192.10.10–14.
5432	TCP	Used by the PostgreSQL database by default. Without an open port in the host firewall, a remote Orchestration Server cannot access the audit database. If you use a different external RDBMS with Orchestration Server, ensure that you open its ports instead.
1443	TCP	Used for Sentinel Syslog secure communications by default, if you configure Sentinel for integration with Cloud Manager.

Figure 2-1 shows the relationships between the Cloud Manager components and managed resources, as well as the traffic flow between them:

Figure 2-1 *Required Network Resources for the Cloud Manager Orchestration Server*

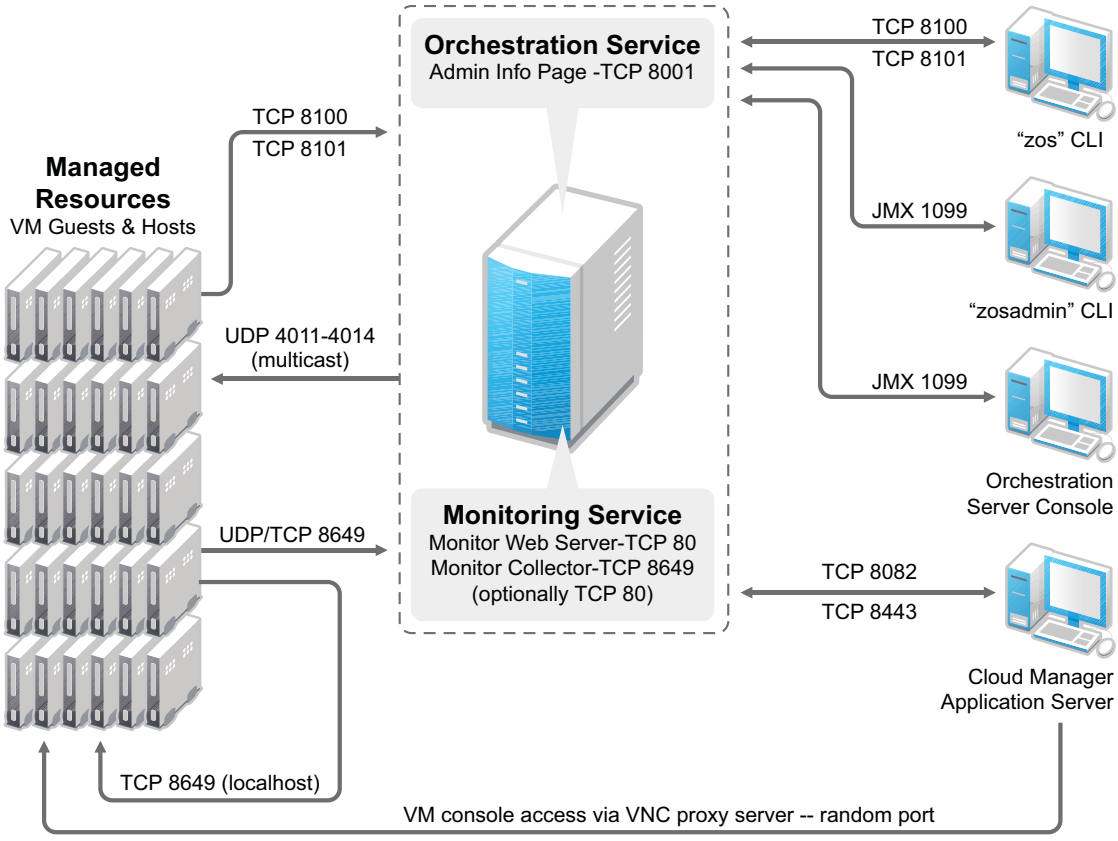

2.3 Cloud Manager Orchestration Agent Requirements

Use the requirements in this section to plan your deployment of the Cloud Manager Orchestration Agent.

- Section 2.3.1, "Requirements for Orchestration Agent," on page 21
- Section 2.3.2, "Port Requirements for the Orchestration Agent," on page 22

2.3.1 Requirements for Orchestration Agent

The physical or virtual machine where you install the Orchestration Agent must meet the requirements in Table 2-3.

Table 2-3 *Orchestration Agent Requirements*

Item	Requirement
Operating System	SUSE Linux Enterprise Server 11 Service Pack 4 (SLES 11 SP 4) on the 64-bit (x86-64) architecture (Intel processor or AMD Opteron processor)

Item	Requirement
Hardware	The Orchestration Agent does not require a minimum hardware configuration other than a minimum recommended disk space of 100 MB.

If you are installing the agent to a vSphere environment, you can install the agent either locally on the vCenter Server (the vCenter appliance is not supported), or on a dedicated system (virtual or physical) as long as its operating system is supported for the Orchestration Agent.

After you install the software on a SLES server, packages for alternative platforms for the Orchestration Agent are available on the server. You can install the agent on most releases of the following operating systems on machines in the same network:

- SUSE Linux Enterprise Server 11 servers
- SUSE Linux Enterprise Desktop 11 desktops
- Red Hat Enterprise Linux 6 servers
- Microsoft Windows Server 2003, 2008, or 2012 servers
- Microsoft Windows Vista, 7, or 8.*x* desktops

For more information, see Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 41.

2.3.2 Port Requirements for the Orchestration Agent

Table 2-4 identifies the ports and protocols required by the Cloud Manager Orchestration Agent. The computing node communicates with the Orchestration Server over a custom protocol. The ports (or their substitutes if you are not using the defaults) must be open for both inbound and outbound communication. Network firewalls must allow outgoing agent connections to these ports on the Orchestration Server.

Table 2-4 *Port Requirements for the Cloud Manager Orchestration Agent*

Port Number	Protocol	Function
8100	TCP	(Not secure) Used with a custom protocol for communication with the Orchestration Server.
8101	TCP	(Secure) Used with a custom protocol for secure communications with the Orchestration Server.

2.4 Cloud Manager Orchestration Console Requirements

The Cloud Manager Orchestration Console is a client software that allows you to administer an Orchestration Server. You can install the console client on the server where you install the Orchestration Server. For more information, see Chapter 2, “Cloud Manager System Requirements,” on page 17.

After you install the Orchestration components on a SLES server, packages for alternative platforms for the Orchestration Console are available on the server. You can install the console on the following operating systems on machines in the same network:

- SUSE Linux Enterprise Server 11 server
- SUSE Linux Enterprise Desktop 11 desktops
- Microsoft Windows Vista, 7, or 8.*x* desktops

For installation information, see Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 48.

2.5 Cloud Manager Application Server Requirements

Use the requirements in this section to plan your deployment of the Cloud Manager Application Server.

- Section 2.5.1, “Requirements for the Application Server,” on page 23
- Section 2.5.2, “Port Requirements for the Application Server,” on page 24

2.5.1 Requirements for the Application Server

The network machine where you install Cloud Manager Application Server software must meet the requirements in Table 2-5:

Table 2-5 *Application Server Requirements*

Item	Requirement
Operating System	SUSE Linux Enterprise Server 11 Service Pack 4 (SLES 11 SP 4) on the 64-bit (x86-64) architecture (Intel processor or AMD Opteron processor)

Hardware	◆ **Processor:**
	  ◆ **Minimum:** 2.5 GHz 64-bit (or equivalent) Intel processor or AMD Opteron processor
	  ◆ **Recommended:** Dual-Core, 2.5 GHz (or greater) 64-bit, Intel processor or AMD Opteron processor
	◆ **RAM:**
	  ◆ **Minimum:** 4 GB
	  ◆ **Recommended:** 6 GB (or greater)
	◆ **Disk Space:** 40 GB
	If the Cloud Manager Application Server and Cloud Manager Orchestration Server are on the same server, the following are minimum requirements:
	◆ 4 Pentium-class CPU cores
	◆ 40 GB disk space
	◆ 4 GB RAM

Item	Requirement
Database	PostgreSQL 9.2 (included with SLES 11 SP 4) You can use a local or external PostgreSQL database for use with the Orchestration Server. We recommend that you install and run PostgreSQL on a different server than where you install Orchestration Server.
Authentication Source	Cloud Manager supports the following authentication methods: ◆ NetIQ eDirectory (LDAP) ◆ Microsoft Active Directory (LDAP) ◆ NetIQ Access Manager

2.5.2 Port Requirements for the Application Server

Table 2-6 identifies the ports and protocols required by the Cloud Manager Application Server. The ports (or their substitutes if you are not using the defaults) must be open for both inbound and outbound communication.

Table 2-6 *Port Requirements for the Cloud Manager Application Server*

Port Number	Protocol	Function
8061	TCP	Used by ESB HTTP.
8102	TCP	(secure) Used by Karaf SSH.
8181	TCP	Used by the Karaf Management console.
8182	TCP	(not secure) Used by Jetty HTTP (default).
8183	TCP	(secure) Used by Jetty HTTPS (default).
10990	TCP	Used by the RMI Registry.
61613	TCP	Used by the active MQ Stomp.
61616	TCP	Used by the active MQ Openwire.
VNC ports (1000–1xxx)	TCP	Used by connections to VM consoles through a VNC client. By default, a VNC proxy port is chosen at random, however the port can be set by the Cloud Administrator in the Configuration page of the Cloud Manager Web Console. There is also an option for an external proxy to offload the traffic from the Cloud Manager Application Server. Typically, this means TCP port 1000 for the first VM on a VM host, 1001 for the second, and so on. These connections go to the VM host, exposing the console on behalf of the VM. For more information, see “Configuring Remote Console Access to Workloads” in the *Cloud Manager Procedures Guide*.

2.6 Cloud Manager Application Console Requirements

The Cloud Manager Application Console is a web-based application. The software runs on the Application Server. For requirements, see Section 2.5, “Cloud Manager Application Server Requirements,” on page 23.

The machine you use to access the Application Console must meet the requirements in Table 2-7:

Table 2-7 *Requirements for the Application Console*

Item	Requirement
Web Browser	Any of the following: ◆ **Microsoft Internet Explorer 9.0 and later:** Supported on Windows 7 (64-bit) and Windows 8.*x* (64-bit) desktop. Ensure that Compatibility View is disabled for the console page. (Open the console page, select **Tools > Compatibility View Settings**, deselect **Display intranet sites in Compatibility View**, then click **Close**.) ◆ **Mozilla Firefox 7.x and later:** Supported on Windows 7 (64-bit) and Windows 8.*x* (64-bit) desktop. ◆ **Apple Safari 5 and later:** Supported on Windows 7 (PC, 64-bit). ◆ **Google Chrome 33 and later:** Supported as a technology preview on Windows 7 (PC, 64-bit) and Windows 8.*x* (64-bit) desktop.
Display Resolution	The minimum requirement is 1024 x 768 with the browser in Full Screen mode (F11).
Pop-Up Blocker	Allow pop-ups from the Cloud Manager Application Server to enable the Help system.

2.7 Cloud Manager Monitoring Agent Requirements

After you install the Cloud Manager Application Server on a SLES server, packages for alternative platforms for the Monitoring Agent are available on the server. You can install the agent on the following operating systems on machines in the same network:

- SUSE Linux Enterprise Server 11 server
- Red Hat Enterprise Linux 6 server
- Microsoft Windows Server 2003, 2008, or 2012 server

For installation information, see Section 5.4, “Alternative Installation Methods for the Cloud Manager Monitoring Agent,” on page 51.

2.8 Requirements and Cloud Manager Support for the Virtual Environment

Cloud Manager supports the virtualization components listed in Table 2-8 for virtualization hypervisors, the host operating system for these hypervisors, and the guest operating systems that run on the virtual machines (VMs) (also called “workloads”). In the Cloud Manager Orchestration Server, the `vsphere` provisioning adapter job is used to provision and manage the lifecycle of the VMs on VMware hosts.

NOTE: Support for VMs is subject to the support of the guest operating system on the target virtualization host by the host vendor. For information about your target VMware hosts, refer to the *VMware Compatibility Guide* (http://www.vmware.com/resources/compatibility/).

For more detail about the life cycle management capabilities of Cloud Manager Orchestration, see Chapter 11, “Configuring the vSphere Provisioning Adapter,” on page 81.

Table 2-8 *VM Technologies with Supported Host Operating Systems, Guest Operating System, and Provisioning Adapter*

Virtualization Hypervisor	Host Operating System	Guest Operating Systems for VMs (Workloads)
VMware ESXi 6.0 U2	Subject to the VMware support matrix	**Microsoft Windows VMs** ◆ Windows Server 2012 R2 (latest SP, 64-bit) ◆ Windows Server 2012 (latest SP, 64-bit) ◆ Windows Server 2008 R2 (latest SP, 64-bit) ◆ Windows Server 2008 (latest SP, 32-bit and 64-bit) ◆ Windows Server 2003 R2 (latest SP, 32-bit only) **Linux VMs** ◆ SUSE Linux Enterprise Server (SLES) 12 (64-bit) ◆ SLES 11 SP3 (64-bit) ◆ SLES 10 SP4 (64-bit) ◆ Red Hat Enterprise Linux (RHEL) 7 ◆ RHEL 6.6 ◆ RHEL 6 (latest SP) ◆ CentOS 7 ◆ CentOS 6.5 ◆ Ubuntu Server 14.04 (32-bit and 64-bit) ◆ Ubuntu Server 13.04 (32-bit and 64-bit) ◆ Ubuntu Server 12.04.3 LTS (32-bit and 64-bit)

Virtualization Hypervisor	Host Operating System	Guest Operating Systems for VMs (Workloads)
VMware ESXi 5.5 (latest update)	Subject to the VMware support matrix	**Microsoft Windows VMs** ◆ Windows Server 2012 R2 (latest SP, 64-bit) ◆ Windows Server 2012 (latest SP, 64-bit) ◆ Windows Server 2008 R2 (latest SP, 64-bit) ◆ Windows Server 2008 (latest SP, 32-bit and 64-bit) ◆ Windows Server 2003 R2 (latest SP, 32-bit only) **Linux VMs** ◆ SUSE Linux Enterprise Server (SLES) 12 (64-bit) ◆ SLES 11 SP3 (64-bit) ◆ SLES 10 SP4 (64-bit) ◆ Red Hat Enterprise Linux (RHEL) 7 ◆ RHEL 6.6 ◆ RHEL 6 (latest SP) ◆ CentOS 7 ◆ CentOS 6.5 ◆ Ubuntu Server 14.04 (32-bit and 64-bit) ◆ Ubuntu Server 13.04 (32-bit and 64-bit) ◆ Ubuntu Server 12.04.3 LTS (32-bit and 64-bit)
VMware ESXi 5.1 (latest update)	Subject to the VMware support matrix	**Windows VMs** ◆ Windows Server 2012 R2 (latest SP, 64-bit) ◆ Windows Server 2012 (latest SP, 64-bit) ◆ Windows Server 2008 R2 (latest SP, 64-bit) ◆ Windows Server 2008 (latest SP, 32-bit and 64-bit) ◆ Windows Server 2003 R2 (latest SP, 32-bit only) **Linux VMs** ◆ SLES 12 (64-bit) ◆ SLES 11 SP3 (64-bit) ◆ SLES 10 SP4 (64-bit) ◆ RHEL 7 ◆ RHEL 6.6 ◆ RHEL 6 (latest SP)[1] ◆ CentOS 7 ◆ CentOS 6.5 ◆ Ubuntu Server 14.04 (32-bit and 64-bit) ◆ Ubuntu Server 13.04 (32-bit and 64-bit) ◆ Ubuntu Server 12.04.3 LTS (32-bit and 64-bit)

[1] For more information about RHEL 6 VM support, see Section 2.8.1, “RHEL 6 VM Support,” on page 28.

2.8.1 RHEL 6 VM Support

You need to be aware of the following limitation of Red Hat Enterprise Linux 6 VMs in the Cloud Manager environment:

- RHEL 6 uses the udev service, which testing has shown renames the network interfaces on a cloned VM and causes configuration errors. To turn off the udev service so that network configuration can work with personalization,
  1. In the file structure of the template VM, open the `/etc/udev/rules.d/70-persistent-net.rules` file and remove all its lines.
  2. In the file structure of the template VM, open the `/lib/udev/write_net_rules` file and comment out (that is, add a # sign preceding the code) the line that looks similar to this:

     ```
     write_rule "$match" "$INTERFACE" "$COMMENT"
     ```

NOTE: Editing the template VM files assures that all its clones will work properly.

2.8.2 Requirements for Machines Designated as VM Hosts

We recommend that computers designated as VM hosts in your data center be able to host the VM and run it according to designated parameters of the specific VM. The processor architecture must match the designated VM’s processor in architecture, although not in version number. In order for a machine to serve as a host machine, it must also have a hypervisor installed along with the operating system.

Table 2-9 *Minimum and Recommended Hardware Requirements for VM Host Machines*

Minimum Requirements	Recommended Hardware
◆ x86_64 ◆ 2 GB RAM ◆ 30 GB hard drive space	◆ x86_64 ◆ 4+ GB RAM ◆ 100+ GB hard drive space

3 Choosing the Installation Patterns and Where to Install Them

Cloud Manager is comprised of a number of different RPM files that are bundled in different installation patterns, all of which are available on the installation media and public key file that you download from NetIQ Downloads (http://dl.netiq.com). Your NetIQ sales representative provides the URL to the media download site, along with the license key you purchased.

NOTE: If you install or configure Cloud Manager components by using a trial key, the product behaves normally for 90 days, although the trial key controls the number of users and managed nodes you can configure. For fully supported functionality, product components require a purchased license key. Contact your NetIQ Sales Representative or a Certified NetIQ Partner for purchase information.

The RPM files in the installation patterns must be installed to a supported version of SUSE Linux Enterprise Server (SLES). The installation uses the **Add-On Products** utility that is available in SUSE's YaST program.

NOTE: After the initial installation and configuration, installers for some Cloud Manager Orchestration components for other operating systems become available in the Orchestration file system.

You can install the Cloud Manager component patterns on machines in your data center according to your own criteria. The information in this section can help you decide which Cloud Manager patterns you want to install and the machines in your data center where you want to install them.

- Section 3.1, "Cloud Manager Orchestration Server Installation Pattern," on page 29
- Section 3.2, "Cloud Manager Installation Pattern," on page 30
- Section 3.3, "Cloud Manager Monitoring Server Installation Pattern," on page 31
- Section 3.4, "Cloud Manager Orchestration Agent Installation Pattern," on page 31
- Section 3.5, "Cloud Manager Monitoring Agent Installation Pattern," on page 32
- Section 3.6, "Cloud Manager Orchestration Console Installation Pattern," on page 32

3.1 Cloud Manager Orchestration Server Installation Pattern

Pattern name [short name]: Cloud Manager Orchestration Server [`zw_zos_server`]

Description: The Orchestration server receives workload instructions from the Cloud Manager Application Server and directs the creation and management of those workloads by the virtual infrastructure. Depending on the size of your virtual infrastructure, you might have one or many Orchestration Servers.

The Orchestration Server requires configuration after installation. To perform the initial configuration, you can use a a text interface at the Linux console (`./config`) or a GUI configuration wizard (`./guiconfig`).

Packages in the pattern: For information about packages that are part of this pattern, mount the product ISO, select the pattern in YaST, and then view its details.

NOTE: Orchestration Server patterns are labeled version 3.5 in the Cloud Manager 2.5 release.

Installation recommendations: Although the machine where you install this server might be capable of handling tasks in addition to the tasks an Orchestration Server performs for Cloud Manager, NetIQ strongly recommends that you install the Orchestration Server software on a dedicated server to ensure optimal performance. For example, you might not want the server to host the Cloud Manager Application Server or NetIQ eDirectory.

NOTE: Although you can install the Orchestration Server on a virtual machine, do not try to manage that VM through the Orchestration Console or other Orchestration Clients.

Installing the server on a VM slows down the performance of the product.

Although it is not mandatory, we recommend that you install and configure the Orchestration Server before you install and configure the Cloud Manager application components.

3.2 Cloud Manager Installation Pattern

Pattern name [short name]: Cloud Manager [`cloudmanager`]

Description: The Cloud Manager Installation Pattern consists of packages for the Cloud Manager Application Server and its Web Console. The Application Server communicates with Cloud Manager Orchestration Servers to provide instructions for provisioning, managing, and removing workloads. It also performs user authentication with a supported authentication source.

The Application Server requires initial configuration after installation to establish authentication with a supported authentication source. The configuration also establishes communication with the Cloud Manager Orchestration Server and its console.

Packages in the pattern: For information about packages that are part of this pattern, mount the product ISO, select the pattern in YaST, and then view its details.

NOTE: You typically install Cloud Manager to use an external PostgreSQL database. However, when you select this pattern, the `postgresql-server` package is selected by default in order to allow you install Cloud Manager to an embedded PostgreSQL server. Deselect the `postgresql-server` package if you plan to use an external PostgreSQL database.

Installation recommendations: Your server might be capable of handling tasks in addition to its Cloud Manager tasks. However, we strongly recommend that you install the Cloud Manager Server software on a dedicated server to ensure optimal performance. For example, you might not want the same server to host the Cloud Manager Orchestration Server or NetIQ eDirectory.

Although not mandatory, we recommend that you install and configure the Orchestration Server before you install and configure the application components.

3.3 Cloud Manager Monitoring Server Installation Pattern

Pattern name [short name]: Cloud Manager Monitoring Server [`zw_mon_server`]

Description: The Cloud Manager Monitoring Server is an Apache Web server that uses open source Ganglia to monitor defined performance data on network resources in a time period you can define.

This server requires configuration after installation. To perform the initial configuration, you can use a a text interface at the Linux console (`./config`) or a GUI configuration wizard (`./guiconfig`).

Packages in the pattern: For information about packages that are part of this pattern, mount the product ISO, select the pattern in YaST, and then view its details.

NOTE: Monitoring Server patterns are all version 3.5 in the Cloud Manager 2.5 release.

Installation recommendations: You can install this server on the same machine with the Orchestration Server, or you can choose any other server with access to the Monitoring Agents.

3.4 Cloud Manager Orchestration Agent Installation Pattern

Pattern name [short name]: Cloud Manager Orchestration Agent [`zw_zos_agent`]

Description: The Cloud Manager Orchestration Agent provides communication between the Orchestration Server and the VM hosts managed by the server. The agent is installed on the VM hosts that run as nodes under the management of the Orchestration Server.

The agent requires configuration after installation. To perform the initial configuration, you can use a text interface at the Linux console (`./config`) or a GUI configuration wizard (`./guiconfig`).

Packages in the pattern: For information about packages that are part of this pattern, mount the product ISO, select the pattern in YaST, and then view its details.

NOTE: Orchestration Agent patterns are labeled version 3.5 in the Cloud Manager 2.5 release.

Installation recommendations: Installing the Cloud Manager Orchestration Agent on the same machine with the Orchestration Server is not supported.

If you are installing the agent to a vSphere environment, you can install the agent either locally on the vCenter Server (the vCenter appliance is not supported), or on a dedicated system (virtual or physical) as long as the OS in that system is supported for the Orchestration Agent.

If you want to support virtual resource management in multiple vSphere environments, we recommend you deploy an Orchestration Agent on a dedicated system. For more information, see "The VMware vSphere Provisioning Adapter" in the *Cloud Manager Administrator Reference*.

You can also install the agent from a Windows installation program or use the Linux pattern files to install to Red Hat Enterprise Linux machines. For more information, see Section 5.2, "Alternative Installation Methods for the Orchestration Agent," on page 41.

3.5 Cloud Manager Monitoring Agent Installation Pattern

Pattern name [short name]: Cloud Manager Monitoring Agent [`zw_mon_agent`]

Description: The Cloud Manager Monitoring Agent can be installed on a server where any other Orchestration pattern is installed, or independently on a SLES or Windows server. The agent installation lays down the Ganglia Agent on each monitored node to collect performance metrics and send the data to the Cloud Manager Monitoring Server.

The agent requires configuration after installation. To perform the initial configuration, you can use a text interface at the Linux console (`./config`) or a GUI configuration wizard (`./guiconfig`).

Packages in the pattern: For information about packages that are part of this pattern, mount the product ISO, select the pattern in YaST, and then view its details.

NOTE: Monitoring Agent patterns are all version 3.5 in the Cloud Manager 2.5 release.

Installation recommendations: If you select the Orchestration Agent pattern, the Monitoring Agent pattern is selected by default. This is only a recommended dependence (most users install both components together) and is not binding. The autoselection is made for your convenience.

Although this agent can be installed using YaST, you can also install it from pattern files located on the ISO image. For more information about these patterns, see Table 5-2, “Monitoring Agent Installation Pattern Files for Linux,” on page 52.

You can also install the agent from a Windows installation program or use the Linux pattern files to install to Red Hat Enterprise Linux machines. For more information, see Section 5.4, “Alternative Installation Methods for the Cloud Manager Monitoring Agent,” on page 51.

3.6 Cloud Manager Orchestration Console Installation Pattern

Pattern name [short name]: Cloud Manager Orchestration Console [`zw_zos_clients`]

Description: The Cloud Manager Orchestration Console is a java-based thick client that administers the functionality of the Orchestration Server from a SLES server or a Windows 7 or 8.*x* desktop on the same network with the Orchestration Server. Before you can perform any Orchestration Server management functions, such as creating user accounts and managing activities of the server, you need to install the Orchestration Console. The console is a thick desktop client designed for administrative tasks including infrastructure management (for example, managing computing resources) and monitoring. You can install the console on the server itself or on another network computer.

This pattern includes both a GUI console and two command line interface tools. These clients let you troubleshoot, initiate, change, or shut down server functions for the Orchestration Server and its computing resources. For information about the client tools, see the *Cloud Manager Administrator Reference*.

Packages in the pattern: For information about packages that are part of this pattern, mount the product ISO, select the pattern in YaST, and then view its details.

NOTE: The Orchestration Console pattern is labeled version 3.5 in the Cloud Manager 2.5 release.

The Orchestration Console and Clients are available as a downloadable Windows installation program (`.exe` file) in the ISO images. For information about using this program, see Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 48.

Installation recommendations: No other Cloud Manager components need to be installed on the machine where you install the console and clients. Provided that the machine where you install the clients can connect with Orchestrate Servers in your data center, where you install the clients is at your discretion.

4 Preinstallation Tasks

Before you install the Cloud Manager components on your anticipated Cloud Manager Orchestration Server or Cloud Manager Application Server, you need to prepare the environment.

- Section 4.1, “Gathering Certificate and License Information,” on page 35
- Section 4.2, “Importing the Public Key for Cloud Manager,” on page 35
- Section 4.3, “Preparing the Server When Multiple NICs and DNS Addresses Exist,” on page 36

4.1 Gathering Certificate and License Information

Before you install Cloud Manager, you need to have the following information available:

- A license key (90-day evaluation license or a full license) is required to use the Cloud Manager Orchestration Server. You should have received this key from NetIQ, then you should have subsequently copied it to the network location that you identify during the pattern installation. Be sure to include the name of the license file in the path.

  If you install or configure Orchestration components by using a trial key, the product behaves normally for 90 days, although the trial key controls the number of users and managed nodes you can configure. For fully supported functionality, product components require a purchased license key. Contact your NetIQ Sales Representative or a Certified NetIQ Partner for purchase information.
- (Optional) Certificate authority information (internal, or signed certificate, private key, and public certificate).

4.2 Importing the Public Key for Cloud Manager

Before you install the Cloud Manager server components, you must import the public key file (such as `netiq-provo-build-key.public.txt`) that you downloaded from the NetIQ Downloads page for Cloud Manager. Import the public key to the server where you will install Cloud Manager Orchestration Server and to the server where you will install Cloud Manager Application Server.

To import the public key:

1. Save the public key file to a location on your anticipated Cloud Manager server.
2. Launch a terminal, then enter one of the following commands:

   ```
   rpm --import <public-key-filename>
   ```

   or

   ```
   gpg --import <public-key-filename>
   ```

4.3 Preparing the Server When Multiple NICs and DNS Addresses Exist

If your anticipated Cloud Manager Orchestration Server has multiple network interfaces and multiple DNS addresses, you need to edit the `/etc/hosts` file on the server to change the default loopback address (127.0.0.1 or 127.0.0.2) to the actual IP address of the server. This is necessary because at server startup, the Orchestration Server tries to determine the `matrix.hostname.full` fact. If the IP address of the host name is found to be a loopback address, it is skipped and subsequently configured incorrectly.

If this change is not made, the **Install Agent** action performed on a VM misconfigures the VM to point to the wrong address (because the grid's `matrix.hostname.full` fact is incorrect), resulting in no connection to the server.

II Installing Cloud Manager

The information in this section provides direction for installing Cloud Manager.

- Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 39
- Chapter 6, “Installing Cloud Manager Application Server Components,” on page 55

5 Installing Cloud Manager Orchestration Components

The RPM files in the Orchestration installation patterns must be installed to a supported version of SUSE Linux Enterprise Server (SLES). After the initial installation and configuration, installers for some Cloud Manager Orchestration components for other operating systems become available in the Orchestration file system.

Some Cloud Manager RPM files have dependencies on SLES patterns that might not have been previously installed on the SLES server. For this reason, we recommend that you mount the SLES install media in a CD ROM drive on the server while you install the Cloud Manager packages, either from another CD ROM drive on the same server, or from a downloaded ISO image.

We recommend that you install and configure the Cloud Manager Orchestration components before you continue with the installation and configuration of Cloud Manager Application Server components. For information about installing Application Server components, see Chapter 6, “Installing Cloud Manager Application Server Components,” on page 55.

NOTE: Testing has shown that installing the Cloud Manager Orchestration Server and the Cloud Manager Application Server on the same machine causes problems with the installation and subsequent operation of Cloud Manager. The product is specifically designed for separate installations of these two complementary servers.

- Section 5.1, “Installing Orchestration Components on SLES,” on page 40
- Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 41
- Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 48
- Section 5.4, “Alternative Installation Methods for the Cloud Manager Monitoring Agent,” on page 51

For information about automated methods you can use to install the Orchestration Agent see Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 41.

For information about uninstalling Cloud Manager components, see Chapter 23, “Uninstalling Orchestration Component Patterns from a SLES Server,” on page 215.

For advanced configuration tasks and methods for optimizing the Cloud Manager Orchestration components, see the *Cloud Manager Administrator Reference*.

5.1 Installing Orchestration Components on SLES

This section describes how to install the Cloud Manager Orchestration components on a server running a supported version of SLES, including the Orchestration Server, Orchestration Agent, the Orchestration Console (accompanied by other Orchestration clients), and the Cloud Manager Monitoring Server and Monitoring Agent.

You should install the Orchestration Server on a dedicated server for optimal performance. You should have already decided which SLES file packages you want to install, and on which machines. If not, the information in Chapter 3, “Choosing the Installation Patterns and Where to Install Them,” on page 29 can help you make that decision.

To install Orchestration components:

1. Log in to the target SLES server as `root`, then open YaST2.
2. Download the appropriate Cloud Manager ISO to the SLES server.

   or

   Load the Cloud Manager DVD on the SLES server.
3. Define the Cloud Manager ISO or DVD as an add-on product:
   - **3a** In the YaST Control Center, click **Software,** then click **Add-On Products**.
   - **3b** Click Add, select **Local ISO Image** or **DVD**, then follow the prompts to add the product.
4. Read and accept the license agreement, then click **Next** to display the Software Selection and System Tasks dialog.

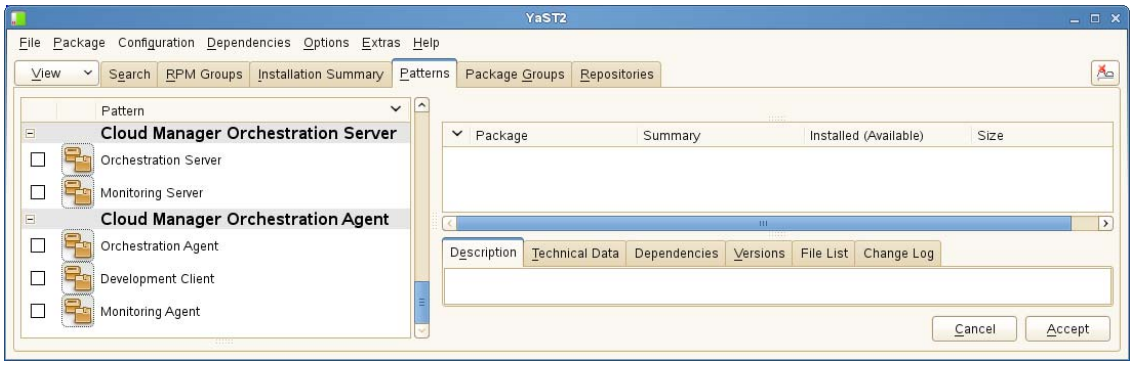

5. Select the installation pattern that contains the Orchestration component packages you want to install on this server.
6. Click **OK** to install the packages.
7. When package installation is complete, click **OK** to close the Installed Add-On Products dialog.

5.2 Alternative Installation Methods for the Orchestration Agent

If you install the Cloud Manager Orchestration Server and the Cloud Manager Application Server, you also need to install the Orchestration Agent on a supported virtual or physical machine so that you can discover resources on such machines and then manage them by using the Cloud Manager Web Console. This section includes information about the installation methods you can use that differ from the standard installation on the Orchestration Server or Application Server.

After you install the software on a SLES server, packages for alternative platforms for the Orchestration Agent are available on the server. You can install the agent on most releases of the following operating systems on machines in the same network:

- SUSE Linux Enterprise Server 11 servers
- SUSE Linux Enterprise Desktop 11 desktops
- Red Hat Enterprise Linux 6 servers
- Microsoft Windows Server 2003, 2008, or 2012 servers
- Microsoft Windows Vista, 7, or 8.*x* desktops

The alternative agent installation methods vary depending on the platform you are installing to. For information about additional requirements for the Orchestration Agent, see Section 2.3, “Cloud Manager Orchestration Agent Requirements,” on page 21.

Agents can be automatically installed on multiple computing resources or groups of computing resources by using your favorite configuration management software. For Windows installation, you can also build your own silent install script. For details about the installation options available for this kind of installation, see Chapter 5.2.5, “Advanced Agent Installation Methods,” on page 45

Windows Installation Source: The Windows installation program for the agent is located on the install media at `\Windows\zosagent_windows_3.5.0_with_jre.exe`. For information about installing the clients on a Windows machine, see Section 5.3.2, “Installing the Console and Clients on Windows,” on page 49.

You can copy this file from the install media to the network, then copy it again to a supported Windows machine where you can run the installation program, or you can open the Administrator Information `.html` page in a web browser. On this page, you can either run the program or download it to copy and run elsewhere. For more information about the Administrator Information page, see Section 5.2.1, “Obtaining the Agent Installer and Supporting Files from the Administrator Information Page,” on page 42.

NOTE: Installation of the Orchestration Agent on a Windows machine does not install the Cloud Manager Monitoring Agent (gmond).

For Monitoring Agent installation information, see Section 5.4.2, “Installing the Cloud Manager Monitoring Agent On Windows Machines,” on page 52.

Linux Installation Source: The manual installation procedure for the agent files on Linux depends on the operating system where you want to install them.

This section includes the following information:

- Section 5.2.1, “Obtaining the Agent Installer and Supporting Files from the Administrator Information Page,” on page 42
- Section 5.2.2, “Installing the Agent on Windows Machines,” on page 43

- Section 5.2.3, “Manually Installing the Agent Packages on SLES Machines,” on page 44
- Section 5.2.4, “Manually Installing the Agent Linux Packages on RHEL Machines,” on page 44
- Section 5.2.5, “Advanced Agent Installation Methods,” on page 45

5.2.1 Obtaining the Agent Installer and Supporting Files from the Administrator Information Page

After you install the Orchestration Server on the network, you can launch the Administrator Information page. The page has links to various installer programs that you can use to install required Cloud Manager software on the computing resources that you will be utilizing in the grid system.

The Orchestration Server Administrator’s web page applications support the following web browsers:

- Microsoft Internet Explorer version 9.0 or later
- Mozilla Firefox version 7.*x* or later
- Google Chrome, version 33 or later

Using a supported browser, enter the following URL to access the Administrator Information from the server:

`http://`*`Orchestration_Server_name`*`:8001/`

This URL is the DNS name (or IP address) of Orchestration Server. Be sure to use Port 8001 in the address to access and display the page, as shown in the following illustration:

Figure 5-1 *Administrator Information Page*

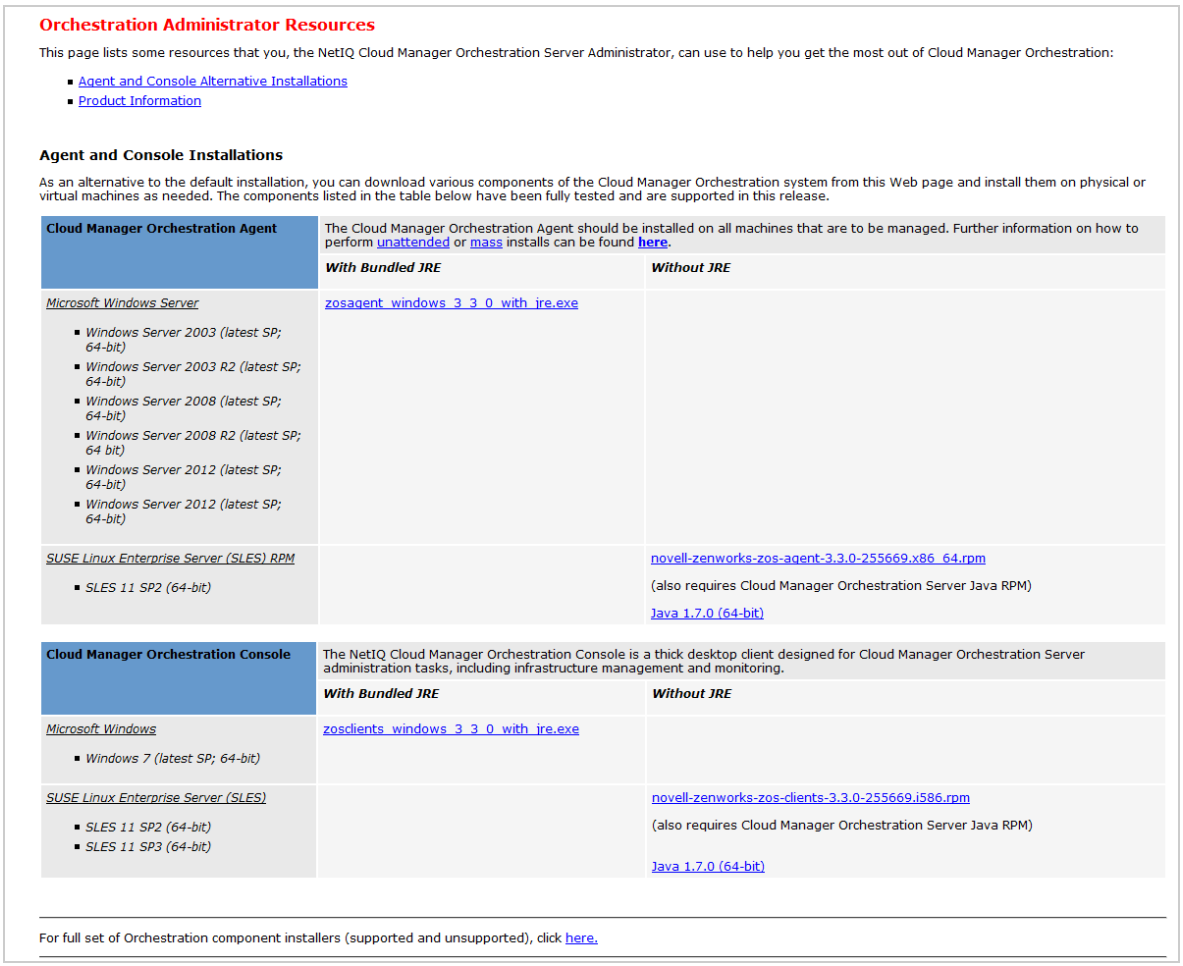

Orchestration Administrator Resources

This page lists some resources that you, the NetIQ Cloud Manager Orchestration Server Administrator, can use to help you get the most out of Cloud Manager Orchestration:

- Agent and Console Alternative Installations
- Product Information

Agent and Console Installations

As an alternative to the default installation, you can download various components of the Cloud Manager Orchestration system from this Web page and install them on physical or virtual machines as needed. The components listed in the table below have been fully tested and are supported in this release.

Cloud Manager Orchestration Agent	The Cloud Manager Orchestration Agent should be installed on all machines that are to be managed. Further information on how to perform unattended or mass installs can be found **here**.	
	With Bundled JRE	***Without JRE***
Microsoft Windows Server ▪ *Windows Server 2003 (latest SP; 64-bit)* ▪ *Windows Server 2003 R2 (latest SP; 64-bit)* ▪ *Windows Server 2008 (latest SP; 64-bit)* ▪ *Windows Server 2008 R2 (latest SP; 64 bit)* ▪ *Windows Server 2012 (latest SP; 64-bit)* ▪ *Windows Server 2012 (latest SP; 64-bit)*	zosagent_windows_3_3_0_with_jre.exe	
SUSE Linux Enterprise Server (SLES) RPM ▪ *SLES 11 SP2 (64-bit)*		novell-zenworks-zos-agent-3.3.0-255669.x86_64.rpm (also requires Cloud Manager Orchestration Server Java RPM) Java 1.7.0 (64-bit)

Cloud Manager Orchestration Console	The NetIQ Cloud Manager Orchestration Console is a thick desktop client designed for Cloud Manager Orchestration Server administration tasks, including infrastructure management and monitoring.	
	With Bundled JRE	***Without JRE***
Microsoft Windows ▪ *Windows 7 (latest SP; 64-bit)*	zosclients_windows_3_3_0_with_jre.exe	
SUSE Linux Enterprise Server (SLES) ▪ *SLES 11 SP2 (64-bit)* ▪ *SLES 11 SP3 (64-bit)*		novell-zenworks-zos-clients-3.3.0-255669.i586.rpm (also requires Cloud Manager Orchestration Server Java RPM) Java 1.7.0 (64-bit)

For full set of Orchestration component installers (supported and unsupported), click here.

The page includes links to information for Cloud Manager Orchestration Server administrators, including product documentation and the installers for the Orchestration Agent.

5.2.2 Installing the Agent on Windows Machines

Cloud Manager requires computing resources in order to run applications. The Orchestration Agent must be installed on each managed device to add that computing resource to the grid where the Orchestration Server can manage it.

To install the agent on a Windows computing resource:

1. Log in to the Windows computer as a user with Administrator privileges.
2. At the location where you copied the Windows agent installer file (`zosagent_windows_<version>_with_jre.exe`), right-click the filename, then select **Run as Administrator**.

   When you launch the installer on Windows Vista, a Security Warning for an Unknown Publisher is displayed. You can ignore this warning and run the installer without a problem.

The welcome page of the Orchestration Agent Setup Wizard is displayed.

3. Click **Next** to display the Select Destination Directory dialog.
4. Accept the default location, then click **Next** to display the Select Start Menu Folder page of the Setup Wizard.
5. Enter the path to the folder where you want the wizard to set up shortcuts to the Agent or select **Next** to accept the default and to display the Windows Services page.
6. Select the services you want to install (at a minimum, you must select **Install Service Orchestration Agent**), then click **Next** to display the Identify Orchestration Server page.
7. Enter the *Orchestration_Server_name* in the **Orchestration Server** field.

   You might find it easier to click **Discover** so that the installer searches for and finds the Orchestration Server on the network. If the installer discovers several servers, ensure that you select the server you previously associated with this agent.
8. Click **Next** to display the Agent Configuration page.

   You can accept the defaults on this page of the Setup Wizard, or you can customize it according to your needs.
9. Click **Next** to run the Orchestration Agent installation until the Agent Setup Wizard completion page is displayed:.
10. Click **Finish** to exit the setup.
11. Register the agent to the Orchestration Console.

    For more information on how to register the agent, see Chapter 10, “Creating a Resource Account,” on page 73.

5.2.3 Manually Installing the Agent Packages on SLES Machines

1. In the Orchestration Agent section, download:
   - Java 1.7.0 (64-bit) (`netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm`)
   - `novell-zenworks-zos-agent-3.5.0-<build_number>.x86_64.rpm`
2. Install the Java 1.7.0 RPM by entering the following command:

   ```
   rpm -ivh netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm
   ```
3. Install the Orchestration Agent by entering the following command:

   ```
   rpm -ivh novell-zenworks-zos-agent-3.5.0-<build_number>.x86_64.rpm
   ```
4. Edit `/opt/novell/zenworks/zos/agent/agent.properties` to set the value of `zos.agent.server` to the IP address of the Orchestration Server where you want to register the agent.
5. Start the agent by entering the following command:

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

5.2.4 Manually Installing the Agent Linux Packages on RHEL Machines

Because you won’t be using the YaST utility to install Orchestration packages on RHEL 6 servers, the information in this section can help you manually install them.

- “Required Agent Installation Files for RHEL Machines” on page 45
- “Manually Installing Orchestration Agents on RHEL 6” on page 45

Required Agent Installation Files for RHEL Machines

The table below lists Orchestration Agent packages that you need to install on RHEL 6 servers. You can find them on the downloaded 64-bit ISO in the `/RHEL6` directory.

Table 5-1 *Required RHEL Installation Packages for the Orchestration Agent*

Platform	Installation Package Name
RHEL 6 (64-bit)	`novell-zenworks-orch-config-3.5.0-<build_number>.noarch.rpm`
	`novell-zenworks-orch-config-gui-3.5.0-<build_number>.noarch.rpm`
	`novell-zenworks-zos-agent-3.5.0-<build_number>.x86_64.rpm` `netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm`

Manually Installing Orchestration Agents on RHEL 6

To install the four packages of the Orchestration Agent on your RHEL 6 server:

1. Download the pertinent 64-bit Add-On ISO from the DVD.
2. Mount the ISO as a loopback device.

   For example, if you are mounting a 64-bit SLES ISO, the command is:

   ```
   $ mount -o loop NetIQ_Cloud_Manager-3.5.0-SLE11.x86_64.iso /mnt
   ```
3. Change your working directory to the location of the RHEL package:

   ```
   $ cd /mnt/RHEL6
   ```
4. Use the package manager included in RHEL to install the Orchestration Agent packages. (Missing dependencies are met by using RHN):

   ```
   $ yum localinstall *.rpm
   ```
5. Run the configuration script:

   ```
   $ /opt/novell/zenworks/orch/bin/config
   ```

   See Section 7.3, “Configuring the Orchestration Agent,” on page 62 for an explanation of the configuration for the Orchestration Agent.

5.2.5 Advanced Agent Installation Methods

This section includes information you can use if you find that the standard and manual methods for installing the Orchestration Agent in your datacenter are inadequate.

- “Silent Installation of the Orchestration Agent” on page 46
- “Using an Orchestration Job to Install the Orchestration Agent on a VM Host” on page 47
- “Automatically Installing the Agent on a VM” on page 48

Silent Installation of the Orchestration Agent

In a large data center, it might not be practical to perform an interactive configuration of the Orchestration Agent on the multiple servers that you intend to use for Cloud Manager resources. The information in this section provides information that can help you perform a silent installation and configuration of the agent.

- “Silent Install and Configuration of the Orchestration Agent for Windows” on page 46
- “Silent Installation and Configuration of the Orchestration Agent RPM” on page 46

Silent Install and Configuration of the Orchestration Agent for Windows

The Cloud Manager Orchestration Server includes an installation help page that provides tips for installing the Windows Orchestration Agent on many machines when you want to use scripting or automation to perform a silent installation.

The page is accessed from the Orchestration Server IP address:

`http://IP_address:8001/install.html`

Figure 5-2 *Orchestration Agent Silent Installation Help*

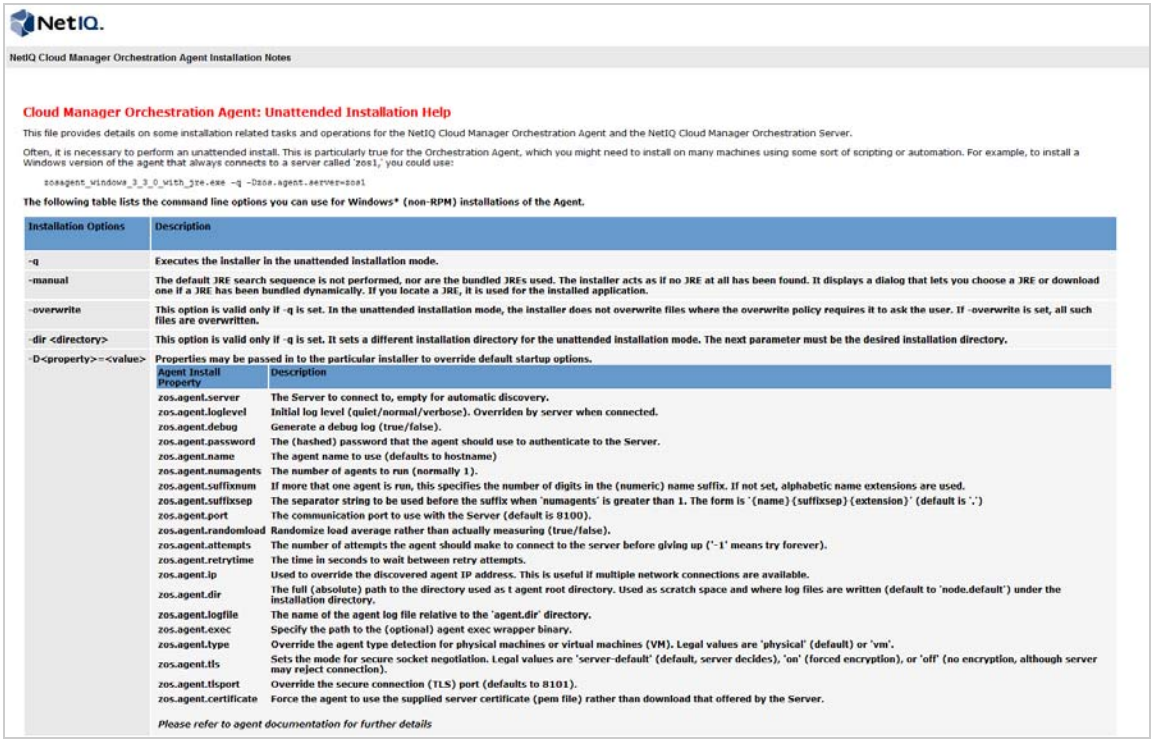
NetIQ.

NetIQ Cloud Manager Orchestration Agent Installation Notes

Cloud Manager Orchestration Agent: Unattended Installation Help

This file provides details on some installation related tasks and operations for the NetIQ Cloud Manager Orchestration Agent and the NetIQ Cloud Manager Orchestration Server.

Often, it is necessary to perform an unattended install. This is particularly true for the Orchestration Agent, which you might need to install on many machines using some sort of scripting or automation. For example, to install a Windows version of the agent that always connects to a server called 'zos1,' you could use:

```
zosagent_windows_3_3_0_with_jre.exe -q -Dzos.agent.server=zos1
```

The following table lists the command line options you can use for Windows* (non-RPM) installations of the Agent.

Installation Options	Description
-q	Executes the installer in the unattended installation mode.
-manual	The default JRE search sequence is not performed, nor are the bundled JREs used. The installer acts as if no JRE at all has been found. It displays a dialog that lets you choose a JRE or download one if a JRE has been bundled dynamically. If you locate a JRE, it is used for the installed application.
-overwrite	This option is valid only if -q is set. In the unattended installation mode, the installer does not overwrite files where the overwrite policy requires it to ask the user. If -overwrite is set, all such files are overwritten.
-dir <directory>	This option is valid only if -q is set. It sets a different installation directory for the unattended installation mode. The next parameter must be the desired installation directory.
-D<property>=<value>	Properties may be passed in to the particular installer to override default startup options.

Agent Install Property	Description
zos.agent.server	The Server to connect to, empty for automatic discovery.
zos.agent.loglevel	Initial log level (quiet/normal/verbose). Overriden by server when connected.
zos.agent.debug	Generate a debug log (true/false).
zos.agent.password	The (hashed) password that the agent should use to authenticate to the Server.
zos.agent.name	The agent name to use (defaults to hostname)
zos.agent.numagents	The number of agents to run (normally 1).
zos.agent.suffixnum	If more that one agent is run, this specifies the number of digits in the (numeric) name suffix. If not set, alphabetic name extensions are used.
zos.agent.suffixsep	The separator string to be used before the suffix when 'numagents' is greater than 1. The form is '{name}{suffixsep}{extension}' (default is '.')
zos.agent.port	The communication port to use with the Server (default is 8100).
zos.agent.randomload	Randomize load average rather than actually measuring (true/false).
zos.agent.attempts	The number of attempts the agent should make to connect to the server before giving up ('-1' means try forever).
zos.agent.retrytime	The time in seconds to wait between retry attempts.
zos.agent.ip	Used to override the discovered agent IP address. This is useful if multiple network connections are available.
zos.agent.dir	The full (absolute) path to the directory used as t agent root directory. Used as scratch space and where log files are written (default to 'node.default') under the installation directory.
zos.agent.logfile	The name of the agent log file relative to the 'agent.dir' directory.
zos.agent.exec	Specify the path to the (optional) agent exec wrapper binary.
zos.agent.type	Override the agent type detection for physical machines or virtual machines (VM). Legal values are 'physical' (default) or 'vm'.
zos.agent.tls	Sets the mode for secure socket negotiation. Legal values are 'server-default' (default, server decides), 'on' (forced encryption), or 'off' (no encryption, although server may reject connection).
zos.agent.tlsport	Override the secure connection (TLS) port (defaults to 8101).
zos.agent.certificate	Force the agent to use the supplied server certificate (pem file) rather than download that offered by the Server.

Please refer to agent documentation for further details

Silent Installation and Configuration of the Orchestration Agent RPM

Use the following process to configure the Orchestration Agent RPM (downloaded from the product ISO) on multiple servers:

1. Perform the product installation and manual configuration of the agent on a “seed” machine. The processes to do this are described in Section 5.1, “Installing Orchestration Components on SLES,” on page 40.
2. On the “seed” machine, copy the file found at `/etc/opt/novell/novell_zenworks_orch_install.conf` to a location where you can modify it locally.

3 Edit the local copy of `novell_zenworks_orch_install.conf`, updating the fields that require a password (for security purposes, when a configuration program runs, the passwords in the `.conf` file are deleted).

4 Edit any other fields as necessary for the configuration of the Orchestration Agent.

5 Distribute the modified file to the machines where you want to perform a silent configuration.

6 At a machine where you distributed the `.conf` file, open YaST and perform the Add-on Installation of the RPM files as described in Section 5.2.3, “Manually Installing the Agent Packages on SLES Machines,” on page 44. Ensure that you do not configure the agent manually.

7 From the bash prompt on the machine where you are configuring the agent, run the following command:

```
/opt/novell/zenworks/orch/bin/config -s -C $CONF_FILE
```

where *CONF_FILE* is the modified configuration file from Step 5.

The silent configuration runs, then the agent is displayed in the Orchestration Console as registered with the server node.

Using an Orchestration Job to Install the Orchestration Agent on a VM Host

The following job code sample shows how you can use a job to install the Orchestration Agent on a VM host.

```
"""
Search for a VM Grid objects using Constraints and run a VM operation on them.
"""
class test(Job):

   def job_started_event(self):

      # collect all VM Instances whose resource ID
      # starts with the string "apache"

      a = AndConstraint()

      e1 = EqConstraint()
      e1.setFact("resource.type")
      e1.setValue("VM")
      a.add(e1)

      e2 = EqConstraint()
      e2.setValue("apache*")
      e2.setMatchMode(EqConstraint.MATCH_MODE_REGEXP)
      e2.setFact("resource.id")
      a.add(e2)

      vms = getMatrix().getGridObjects(TYPE_RESOURCE,a,None)
      for vm in vms:
         vm.installAgent()
```

Automatically Installing the Agent on a VM

To automatically install the Orchestration Agent on a VM that you created in the client, right-click a VM that has been shut down, then select **Install Agent**. This launches a job that installs the Orchestration Agent on the VM, regardless of its platform. The agent's service is started the next time you provision the VM.

5.3 Alternative Installation Methods for the Orchestration Console and Clients

You can use the Cloud Manager Orchestration Console to administer the Orchestration Server from the server, or separately from other machines on the same network. This section includes information about the installation methods you can use that differ from the standard installation on the Orchestration Server.

After you install the Orchestration components on a SLES server, packages for alternative platforms for the Orchestration Console are available on the server. You can install the console on the following operating systems on machines in the same network:

- SUSE Linux Enterprise Server 11 server
- SUSE Linux Enterprise Desktop 11 desktops
- Microsoft Windows Vista, 7, or 8.*x* desktops

The alternative console installation methods vary depending on the platform you are installing to. For information about additional requirements for the Orchestration Console, see Section 2.4, "Cloud Manager Orchestration Console Requirements," on page 22.

Windows Installation Source: The Windows installation program for the console and clients is located on the install media at `\Windows\zosclients_windows_3_2_0_with_jre.exe`. For information about installing the clients on a Windows machine, see Section 5.3.2, "Installing the Console and Clients on Windows," on page 49.

You can copy this file from the install media to the network, then copy it again to a supported Windows machine where you can run the installation program, or you can open the Administrator Information `.html` page in a web browser. On this page, you can either run the program or download it to copy and run elsewhere.

Linux Installation Source: The manual installation procedure for the client files on Linux depends on the operating system where you want to install them. For information about installing the clients on a Linux machine, see Section 5.3.3, "Installing the Console and Clients on a SLES Server," on page 51.

- Section 5.3.1, "Obtaining Installers from the Administrator Information Page," on page 49
- Section 5.3.2, "Installing the Console and Clients on Windows," on page 49
- Section 5.3.3, "Installing the Console and Clients on a SLES Server," on page 51

5.3.1 Obtaining Installers from the Administrator Information Page

After you install the Orchestration Server on the network, you can launch the Administrator Information page. The page has links to various installer programs that you can use to install required Cloud Manager software on the computing resources that you will be utilizing in the grid system.

The following browsers support the Orchestration Server Administrator's web page applications:

- Microsoft Internet Explorer version 9.0 or later
- Mozilla Firefox version 7.*x* or later
- Google Chrome, version 33 or later

Using a supported browser, enter the following URL to access the Administrator Information from the server:

`http://`*`Orchestration_Server_name`*`:8001/`

This URL is the DNS name (or IP address) of Orchestration Server. Be sure to use Port 8001 in the address to access and display the page, as shown in the following illustration:

Figure 5-3 *Administrator Information Page*

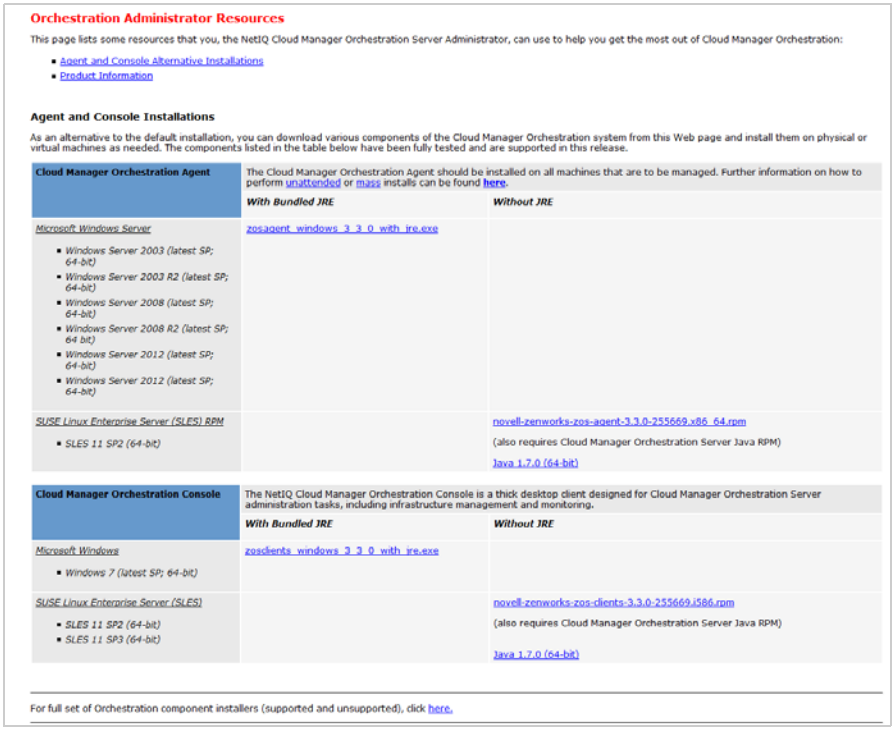

The page includes links to Orchestration information for data center administrators, including links to product documentation and to the installers for the Orchestration Console and clients

5.3.2 Installing the Console and Clients on Windows

1 At a Windows desktop location where you downloaded the client installer file for Windows, double-click the `zosclients_windows_3_3_0_with_jre.exe` icon to run the installer.

When you launch the installer, a Security Warning for an Unknown Publisher is displayed. You can ignore this warning and run the installer without a problem.

The first page of the Cloud Manager Orchestration Tools Setup Wizard is displayed.

2 Click **Next** to display the License Agreement page.

3 Accept the license agreement, then click **Next** to display the Select Destination Directory page.

4 Select the folder where you want to install the clients, then click **Next** to display the Select Start Menu folder page.

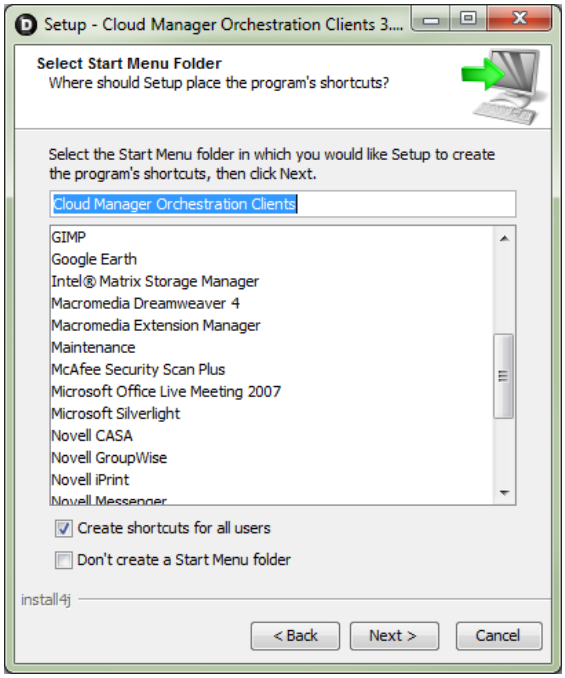

5 Select the Start Menu folder where you want the install program to create the client shortcuts, then click **Next** to begin the installation. The file extraction and copy process proceeds until the Cloud Manager Orchestration Tools Setup Completion page is displayed.

The following items are installed on the Windows machine:

- **Custom Orchestration Tools:** This includes the zos command line tool and a `.jar` file used to develop custom clients. The zos command line tool provides a non-web method for a user to access the server.

  For more information, see the *Cloud Manager Administrator Reference*.

- **Orchestration Console and Command Line:** This includes the Cloud Manager Orchestration Console, which is a thick client console for administrators. It also installs the zosadmin command line tool for administrators. Both of these tools require administrator login.

  For more information, see the *Cloud Manager Administrator Reference*.

6 Click **Finish** to exit the setup.

Installing these components on a Windows workstation adds several items to the program group available from **Start > All Programs > Novell > ZOS > Clients**. One of these programs is the Orchestration Server Command Prompt. The PATH is preset in this prompt to run the zos and zosadmin commands.

5.3.3 Installing the Console and Clients on a SLES Server

1. In the Orchestration Clients section of the Administrator Information page, download:
   - Java 1.7.0 (64-bit) (`netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm`)
   - `novell-zenworks-zos-clients-3.5.0-<build_number>.<platform>.rpm`
2. Install the Java 1.7.0 RPM by entering the following command:

   ```
   rpm -ivh netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm
   ```
3. Install the Orchestration Console by entering the following command:

   ```
   rpm -ivh novell-zenworks-zos-clients-3.5.0-<build_number>.<platform>.rpm
   ```

5.4 Alternative Installation Methods for the Cloud Manager Monitoring Agent

The Cloud Manager Monitoring Agent can be installed on a server where any other Orchestration pattern is installed, or independently on another machine in the same network. This section includes information about the installation methods you can use that differ from the standard installation.

After you install the Cloud Manager Application Server on a SLES server, packages for alternative platforms for the Monitoring Agent are available on the server. You can install the agent on the following operating systems on machines in the same network:

- SUSE Linux Enterprise Server 11 server
- Red Hat Enterprise Linux 6 server
- Microsoft Windows Server 2003, 2008, or 2012 server

The alternative agent installation methods vary depending on the platform you are installing to. For information about additional requirements for the Monitoring Agent, see Section 2.7, “Cloud Manager Monitoring Agent Requirements,” on page 25.

The agent installation lays down the Ganglia Agent on each monitored node to collect performance metrics and send the data to the Cloud Manager Monitoring Server.

If you need to install Cloud Manager Monitoring Agent files without using the standard SLES installation script, the information in this section can help you identify the installation source files and give you some direction on how to install them to Linux or Windows machines.

- Section 5.4.1, “Installing the Cloud Manager Monitoring Agent on Linux Servers,” on page 52
- Section 5.4.2, “Installing the Cloud Manager Monitoring Agent On Windows Machines,” on page 52

For more information about agent installation, see Appendix 5.2.5, “Advanced Agent Installation Methods,” on page 45.

5.4.1 Installing the Cloud Manager Monitoring Agent on Linux Servers

This section can help you identify the correct installation files for the Cloud Manager Monitoring Agent on the Cloud Manager product ISO and provide you with some installation instructions for installing those files on SLES or RHEL servers.

- “Cloud Manager Monitoring Agent Installation Files for Linux Servers” on page 52

Cloud Manager Monitoring Agent Installation Files for Linux Servers

The Cloud Manager Monitoring Agent requires that you install the agent program file, appropriate for the operating system:

Table 5-2 *Monitoring Agent Installation Pattern Files for Linux*

Operating System	Installation File

SLES 11 SP 4 (64-bit)	`<cd>/suse/setup/descr/zw_mon_agent-<cmos_version>-0.x86_64.pat`
	To install this pattern file with the `zypper` command, enter
	`zypper in -t pattern zw_mon_agent-<cmos_version>-0.x86_64.pat`

RHEL 7 (latest update, 64-bit)	`<cd>/RHEL7/novell-zenworks-monitor-gmond-<version>-67.1.x86_64.rpm`
RHEL 6 (latest update, 64-bit)	`<cd>/RHEL6/novell-zenworks-monitor-gmond-<version>-67.1.x86_64.rpm`

5.4.2 Installing the Cloud Manager Monitoring Agent On Windows Machines

Installing the Cloud Manager Orchestration Agent on Microsoft Windows Server 2003 or Windows Server 2008 (see Section 5.2.2, “Installing the Agent on Windows Machines,” on page 43) does not automatically install the Cloud Manager Monitoring Agent.

A separate installation package is available for installing the Monitoring Agent on Windows platforms where you have installed the Orchestration Agent.

Table 5-3 *Monitoring Agent Installation Files for Windows*

Operating System	Installation File
Windows Server 2012 R2 (64-bit Windows Server 2008 R2 latest SP (with Hyper-V role, 64-bit) Windows Server 2008 R2 latest SP (64-bit) Windows Server 2003 R2 latest SP (64-bit) Windows Server 2003 latest SP (64-bit)	`<cd>/Windows/GmondSetup.exe`

This section includes the following information:

- “Installing the Cloud Manager Monitoring Agent for Windows” on page 53
- “Configuring the Monitoring Agent for Windows” on page 54

Installing the Cloud Manager Monitoring Agent for Windows

- “Hardware and Software Requirements” on page 53
- “Installing the Monitoring Agent” on page 53
- “Starting and Stopping the Monitoring Agent” on page 53
- “Uninstalling the Monitoring Agent” on page 53

Hardware and Software Requirements

The Cloud Manager Monitoring Agent (gmond) can be installed only on a machine running a supported version of Windows Server where the Orchestration Agent is also installed. Only a 32-bit version of gmond is provided, but it will run normally on 64-bit systems. It requires the same minimum hardware configuration as the Orchestration Agent. Port 8649 must be available for gmond to communicate with the agent.

Installing the Monitoring Agent

Use these steps to install the Monitoring Agent as a service on a Windows Server machine.

1. Download the pertinent Add-On ISO from the product DVD.
2. Create a CD from the ISO or use *ISOBuster* (or a similar tool) to mount the ISO.
3. Browse to the `Windows/` folder and search for the Cloud Manager Monitoring Agent.
4. Double-click the Monitoring Agent icon (`gmondsetup.exe`) and follow the wizard through the setup:

NOTE: You must be logged on as Administrator to run the installation program. If you are installing on Windows Server 2008 or later, click **Accept** in the User Account Control dialog to allow the installation to proceed as the Windows Administrator user.

Starting and Stopping the Monitoring Agent

If you are logged on as the Windows administrator, you can start and stop the gmond service by using the Windows Services Control Panel.

1. From the desktop, click **My Computer**, select **Manage**, expand **Services and Applications**, expand **Services**, then right-click the gmond service object and choose the **Start** or **Stop** option as needed.

Uninstalling the Monitoring Agent

You can uninstall gmond by using the Add/Remove Programs utility in the Windows Control Panel. Uninstalling gmond automatically shuts down the gmond service prior to uninstalling.

Configuring the Monitoring Agent for Windows

This version of the Monitoring Agent uses the open source gmond 3.1.7 code. It is preconfigured to run only on a Windows local machine. It does not support multicasting. The installation includes a preconfigured `gmond.conf` that works only on the local host. You can manually edit `gmond.conf` for a different configuration, but your changes are not supported by NetIQ.

If you choose to customize an unsupported configuration for your needs, you can test the configuration by stopping the gmond service and then restarting gmond from the Windows command line prompt. Use the `-d` (debug) and `-f` (foreground) options to capture any error messages generated by the new configuration.

6 Installing Cloud Manager Application Server Components

Cloud Manager transforms your virtual infrastructure into a true Cloud environment. Built to operate with your existing VMware virtual hosts, Cloud Manager accelerates delivery of services through on-demand requesting of workloads and automated provisioning of the workloads.

Some Cloud Manager RPM files have dependencies on SUSE Linux Enterprise Server (SLES) patterns that might not have been previously installed on the SLES server. For this reason, we recommend that you mount the SLES install media in a disk drive on the server while you install the Cloud Manager packages, either from another disk drive on the same server or from a downloaded ISO image.

We recommend that you install and configure the Cloud Manager Orchestration components before you continue with the installation and configuration of Cloud Manager Application Server components. For information about installing Orchestration components, see the Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 39.

NOTE: Testing has shown that installing the Cloud Manager Orchestration Server and the Cloud Manager Application Server on the same machine causes problems with the installation and subsequent operation of Cloud Manager. The product is specifically designed for separate installations of these two complementary servers.

This section includes information about installing the Cloud Manager RPM files to a server in your data center.

- Section 6.1, “Installing Cloud Manager Pattern on Your SLES Server,” on page 55

6.1 Installing Cloud Manager Pattern on Your SLES Server

The RPM files in the Cloud Manager installation patterns must be installed to a supported version of SUSE Linux Enterprise Server (SLES). You should install the Application Server on a dedicated server for optimal performance. For more information about the installation requirements for Cloud Manager, see Chapter 2, “Cloud Manager System Requirements,” on page 17.

1. Log in to the target SLES server as `root`, then open YaST.
2. From the NetIQ product downloads website, download the appropriate Cloud Manager ISO to the SLES server.

   or

   Load the Cloud Manager DVD on the SLES server.
3. Define the Cloud Manager ISO or DVD as an add-on product:
   - **3a** In the YaST Control Center, click **Software,** then click **Add-On Products**.
   - **3b** Click **Add**, select **Local ISO Image** or **DVD**, then follow the prompts to add the product.
4. Read and accept the license agreement, then click **Next** to display the Software Selection and System Tasks dialog.

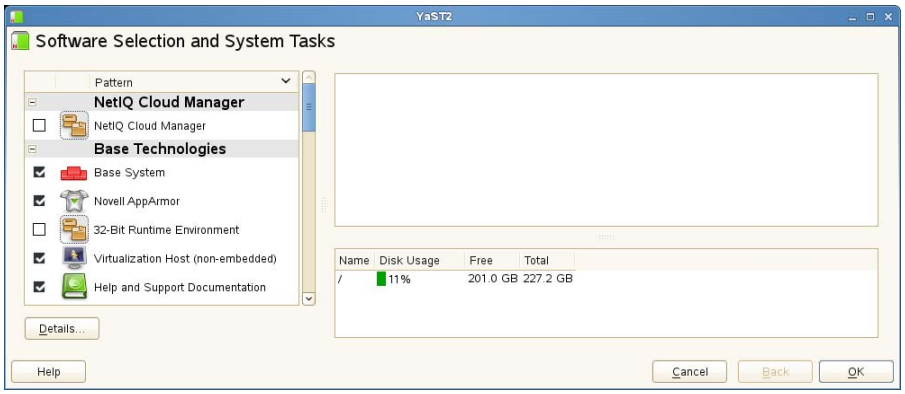

5 Select the **Cloud Manager** installation pattern.

6 Click **OK** to install the packages.

7 When package installation is complete, click **OK** to close the Installed Add-On Products dialog.

III Configuring Cloud Manager

This part includes information you need as you install the Cloud Manager components:

- Chapter 7, “Configuring Cloud Manager Orchestration Components,” on page 59
- Chapter 8, “Configuring Connections to the Cloud Manager Application Server,” on page 67
- Chapter 9, “Launching the Orchestration Console and Logging in to the Orchestration Server,” on page 71
- Chapter 10, “Creating a Resource Account,” on page 73
- Chapter 11, “Configuring the vSphere Provisioning Adapter,” on page 81
- Chapter 12, “Configuring vSphereUpdate Client and Monitoring Jobs,” on page 95
- Chapter 13, “Configuring Sysprep or Autoprep,” on page 101
- Chapter 14, “Using the Cloud Manager Application Server Configuration Tool,” on page 121

7 Configuring Cloud Manager Orchestration Components

This section discusses the basic configuration of all Cloud Manager Orchestration components after each is installed. Component configuration is done either with a text-based configuration tool or with a GUI Wizard configuration tool.

The text-based configuration script detects which RPM patterns are installed, but the GUI Configuration Wizard requires that you specify the components to be configured, whether the patterns have been installed on the server or not.

It is possible to execute the text-based configuration file Orchestration components from the Cloud Manager configuration utility, but this occurs only if you install Cloud Manager Application components on the same server as the Cloud Manager Orchestration components, which is only likely if you are setting up your system for a demonstration.

Both the text-based tool and the GUI Wizard tool produce a configuration file that can be used to automatically reconfigure your system after an upgrade. If you use the tools to reconfigure your server after the original configuration has been done, ensure that you reconfigure all of the components that are installed on the system (this is the default).

NOTE: Remember that the Cloud Manager Orchestration components are version 3.5. This might cause confusion because they are to be used with Cloud Manager Application components, which are version 2.5.

Some Considerations When Configuring with the GUI Wizard

If you have only a keyboard to navigate through the pages of the GUI Configuration Wizard, use the Tab key to shift the focus to a control you want to use (for example, a **Next** button), then press the Spacebar to activate this control.

When you have finished answering the configuration questions in the wizard, the Cloud Manager Orchestration Configuration Summary page displays. Although this page of the wizard lets you navigate by using the Tab key and the Spacebar, you need to use the Ctrl+Tab combination to navigate past the summary list. Click **Back** if you accidentally enter the summary list, and re-enter the page to navigate to the control buttons.

By default, the **Configure now** check box on the page is selected. If you accept this default, the wizard starts the Orchestration Server and applies the configuration settings. If you deselect the check box, the wizard writes out the configuration file to `/etc/opt/novell/novell_zenworks_orch_install.conf` without starting the Orchestration Server or applying the configuration settings.

You can use this `.conf` file to start the Orchestration Server or Agent and apply the settings either manually or with an installation script. Use the following command to run the configuration:

```
/opt/novell/zenworks/orch/bin/config -rs
```

This section includes the following information:

- Section 7.1, “Configuring the Orchestration Server,” on page 60
- Section 7.2, “Configuring the Monitoring Server and Monitoring Agent,” on page 62
- Section 7.3, “Configuring the Orchestration Agent,” on page 62
- Section 7.4, “Validating and Optimizing the Orchestration Installation,” on page 64
- Section 7.5, “Orchestration Server Connection and Read Timeout Values,” on page 64

When the installation and configuration are complete, you need to validate and optimize the configuration. You can them proceed to register the resources to be managed by the Cloud Manager system. Refer to Chapter 10, “Creating a Resource Account,” on page 73 for detailed information about getting resources to manage in the Cloud Manager system.

7.1 Configuring the Orchestration Server

Because so much of Cloud Manager’s operations depends on the Orchestration Server, we recommend that you configure it before you configure any other Cloud Manager component.

1 Ensure that you are ready with the information that you will be prompted for during the configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Configuration Type	Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation. This section discusses standard installation, so specify `s` (for `standard`) or press Enter to accept the default. For more information about High Availability configuration, see the Chapter 15, “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support,” on page 135.
Server IP Address	You need to know the IP address of this server if you have multiple network interfaces. If specified, the server uses the specified host name or IP address for binding RMI connections.
Grid Name	A grid is an administrative domain container holding all of the objects in your network or data center. The Orchestration Server monitors and manages these objects, including users, resources, and jobs. The grid name you create here is displayed as the name for the container placed at the root of the tree in the Explorer panel of the Orchestration Console.
Administrator User	The name you specify here is required when you access the Orchestration Console or the `zosadmin` command line interface. You should remember this user name for future logins.
Administrator Password	The password you specify here is required when you access the Orchestration Console or the `zosadmin` command line interface. You should remember this user name for future logins.

Server Configuration Requirement	Explanation and Action
Auditing Database JDBC URL	If you answer yes to this question, you need access to a relational database management system. We recommend that this database be installed on a different server from where you installed Cloud Manager. NetIQ has tested and supports only the PostgreSQL relational database as the audit database for this release of Cloud Manager. If you use a different RDBMS, no support or documentation is available from NetIQ. For more information, see Chapter 16, “Configuring the Orchestration Server to Use an Audit Database,” on page 149.
Path to License File	A license key (90-day evaluation license or a full license) is required to use this product. You should have received this key from NetIQ, then you should have subsequently copied it to the network location that you specify here. Be sure to include the name of the license file in the path.
User Portal	This utility is no longer supported. Select the default and continue.
Admin Info Port	Port 8001 on the Orchestration Server provides access to an Administrator Information page that includes links to product documentation, agent and client installers, and product tools to help you understand and use the product. Specify another port number if 8001 is reserved for another use on this server.
Orchestration Agent Port	Port 8100 is used for communication between the Orchestration Server and the Orchestration Agent. Specify another port number if 8100 is reserved for another use. If your Orchestration Server communicates with ESX servers, we recommend you configure port 8101. This requires that you configure all other Orchestration Agents communicating with this server to use port 8101. This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the **Configure Advanced Settings** check box in the wizard, you have the option of changing the default values. If you leave the check box deselected the setting is configured with normal defaults.
(Optional) Path to TLS Server Certificate and TLS Server Private Key	A PEM-encoded TLS certificate and key is needed for secure communication between the Orchestration Server and Orchestration Agent. If you do not want the Orchestration Server to generate a certificate and key for authentication, you need to provide the location of an existing certificate and key. This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the **Configure Advanced Settings** check box in the wizard, this parameter is listed, but default values are provided only if the previous value is manually set to **no**.

2 At the computer where you installed the Cloud Manager Orchestration Server pattern, run the Cloud Manager Orchestration configuration utility:

`/opt/novell/zenworks/orch/bin/config`

or

`/opt/novell/zenworks/orch/bin/guiconfig`

3 Follow the prompts to complete the configuration.

7.2 Configuring the Monitoring Server and Monitoring Agent

The Cloud Manager Monitoring Server leverages open source (Ganglia) monitoring of the performance of certain data on network resources in a time period you define. The network resources being monitored must have the Cloud Manager Monitoring Agent installed.

1 Ensure that you are ready with the information that you are prompted for during the Monitoring Server configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Is this computer to be a Monitoring Server or a Monitored Node?	If you chose to install the Monitoring Server pattern, the Monitoring Agent pattern is installed on the same computer by default. You can also install the Monitoring Agent pattern independent of the Monitoring Server, but you should install it on the same node where you install an Orchestration Agent. The configuration lets you choose not to configure this computer as a Monitoring Server.
Hostname or IP Address of the Monitoring Server	You need to know the host name or IP address of the server if you configured as the Cloud Manager Monitoring Server. Monitored nodes send their metrics to this address.
Monitored Computer Name	The descriptive name you designate appears in the monitoring interface as the name or location of the monitored node.

2 At the computer where you installed the Cloud Manager Monitoring Server or Cloud Manager Monitoring Agent pattern, run the configuration utility:

`/opt/novell/zenworks/orch/bin/config`

or

`/opt/novell/zenworks/orch/bin/guiconfig`

3 Follow the prompts to complete the configuration of the Monitoring Server or the Monitoring Agent.

7.3 Configuring the Orchestration Agent

The Cloud Manager Orchestration Agent manages the life cycle of VMs in your hypervisor environment under the direction of the Orchestration Server. You install the agent on computers where those VMs reside.

1 Ensure that you are ready with the information that you will be prompted for during the Monitoring Server configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Configuration Type	Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation. This section discusses standard installation, so specify `s` (for `standard`) or press Enter to accept the default. For more information about High Availability configuration, see the Chapter 15, “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support,” on page 135.
Agent Name	The Orchestration Agent requires a name to authenticate to the Orchestration Server.
Orchestration Server Hostname or IP Address	The DNS name or IP address of the Orchestration Server that this agent binds to.
Always Implement the Orchestration Server Certificate and Key?	The Agent relies on the Orchestration Server’s TLS certificate as verification that it is communicating with the correct Orchestration Server. Decide whether you want to always trust the server certificate after the agent initially downloads it from the server, or if you want to exercise the certificate and key every time the agent connects to the server.
Is the Node a Physical or a Virtual Machine?	If the computer where you installed the agent is actually a VM, the Cloud Manager Server approaches its management in a unique way.
Agent Port	Port 8100 is used for communication between the Orchestration Server and the Orchestration Agent. Specify another port number if 8100 is reserved for another use. For an Agent installed on ESX, configure port 8101.
(Optional) Agent IP Address	You can specify a local bind address for the agent if you want to. If you specify an address, the agent tries to use this address locally when it connects to the Orchestration Server. If you don’t specify an address, the operating system automatically sets the local address for each connection.
Path to Server Certificate	Specify the path to the Orchestration Server certificate file. The default path is `/root/zos_server_cert.pem`. **NOTE:** This configuration parameter is considered an advanced setting for the Orchestration Agent in the GUI Configuration Wizard, but only if you set **Provide Existing Orchestration Server Certificate** to **yes**.

2 At the computer where you installed the Cloud Manager Orchestration Agent pattern, run the configuration utility:

```
/opt/novell/zenworks/orch/bin/config
```

or

```
/opt/novell/zenworks/orch/bin/guiconfig
```

3 Follow the prompts to complete the configuration of the Orchestration Agent.

NOTE: Some configuration parameters for the agent are considered “advanced setting” in the GUI Configuration Wizard. If you select the **Configure Advanced Settings** check box in the wizard, the setting is configured with normal defaults. Leaving the check box deselected lets you have the option of changing the default value.

You can also configure the agent as part of the silent installation procedure. See “Silent Installation of the Orchestration Agent” on page 46.

7.4 Validating and Optimizing the Orchestration Installation

1. Open the configuration log file (`/var/opt/novell/novell_zenworks_orch_install.log`) to ensure that the components were correctly configured.
2. Access the Administrator Information Page to verify that the Orchestration Server is installed and running. Use the following URL to open the page in a web browser:

   `http://`*`DNS_name_or_IP_address_of_Orchestration_Server`*`:8001`

   The Administrator Information page includes links to separate installation programs (installers) for the Orchestration Agent and the Orchestration Clients. The installers are used for various operating systems. You can download the installers and install the agent or the clients on any supported machine you choose. For more information, see Section 5.2.3, “Manually Installing the Agent Packages on SLES Machines,” on page 44.

If you installed the Orchestration Tools, you can increase the heap size that the JVM handles. This enables the console to manage a larger number of objects.

1. Open the `zoc` bash shell script at `/opt/novell/zenworks/zos/server/bin`.

   On Microsoft Windows, the path to the console is `program files\novell\zos\clients\bin\zoc.bat`. For more information, see Section 5.2.3, “Manually Installing the Agent Packages on SLES Machines,” on page 44.
2. Inside the script, find the following line where the JVM parameters are defined:

   ```
   JVMARGS="-Xmx256m -Xms256m -Xmn64m -XX:NewSize=64m -XX:MaxNewSize=64m"
   ```

   The `-Xmx` argument specifies the maximum heap size for the JVM. Increasing the heap size prevents a JVM out of memory condition.
3. Change the value in the `-Xmx` argument from 256MB to 512MB.

If you want to reconfigure the components of a Cloud Manager Orchestration system that you previously installed and configured, you can rerun the configuration script or the GUI Configuration Wizard and change your responses during the configuration process.

7.5 Orchestration Server Connection and Read Timeout Values

Cloud Manager 2.5 adds reliability for zone checking by adding the ability to specify a connection timeout value and a read timeout value in the `/etc/com.novell.cm.core.service.cfg` file on your Orchestration Server. The default timeout values are 30 seconds for Connection and 60 seconds for Read.

To edit the timeout values:

1. On the Orchestration Server, open the `/etc/com.novell.cm.core.service.cfg` file in a text editor.
2. In the file, modify the appropriate values in seconds for the timeout parameters:

   ```
   psoConnectionTimeoutSeconds=30
   psoReadTimeoutSeconds=60
   ```

3. Save the file.
4. Restart the Orchestration Server to apply the modified settings.

8 Configuring Connections to the Cloud Manager Application Server

The Cloud Manager Application Server requires an HTTP connection to each Cloud Manager Orchestration Server that you want to define as a zone in your Cloud Manager system. This connection can be secure (SSL) or non-secure (no SSL).

The following sections provide instructions for enabling secure and non-secure connections. You must complete the instructions for each Cloud Manager Orchestration Server that you plan to define as a Cloud Manager zone.

- Section 8.1, “Enabling a Secure Connection,” on page 67
- Section 8.2, “Enabling a Non-Secure Connection,” on page 69

8.1 Enabling a Secure Connection

A secure connection requires certificate authentication between the Cloud Manager Application Server and the Cloud Manager Orchestration Server.

The first time it is started, the Cloud Manager Web Service creates a keystore, generates a public/private key pair, and exports the public key to a certificate. The Web Service is started automatically as part of the Cloud Manager Orchestration Server startup or manually by using the following command:

```
/etc/init.d/novell-pso-ws start
```

To complete the configuration of the secure connection, you need to import the Cloud Manager Web Service’s public certificate to the Cloud Manager Application Server trust store and configure the secure port for the Cloud Manager Web Service. The following sections provide instructions:

- Section 8.1.1, “Configuring the Cloud Manager Web Service Secure Port,” on page 67
- Section 8.1.2, “Disabling SSL v3 in Java Jetty,” on page 68

8.1.1 Configuring the Cloud Manager Web Service Secure Port

By default, the Cloud Manager Web Service listens on port 8443. You can change this port if necessary.

1. On the Cloud Manager Orchestration Server, open the `jetty-ssl.xml` file:

   ```
   /etc/opt/novell/pso-ws/jetty/jetty-ssl.xml
   ```

2. Locate the `<Call name ="addConnector">` section. It will look similar to the section shown below:

```
<Call name="addConnector">
  <Arg>
    <New class="org.mortbay.jetty.security.SslSocketConnector">
      <Set name="Port">8443</Set>
      <Set name="maxIdleTime">30000</Set>
      <Set name="handshakeTimeout">2000</Set>
      <Set name="keystore"><SystemProperty name="jetty.home" default="."
        />/etc/keystore</Set>
      <Set name="password">OBF:1vny1zlo1x8e1vnw1vn61x8g1zlu1vn4</Set>
      <Set name="keyPassword">OBF:1u2u1wml1z7s1z7a1wnl1u2g</Set>
      <Set name="truststore"><SystemProperty name="jetty.home"
        default="." />/etc/keystore</Set>
      <Set name="trustPassword">OBF:1vny1zlo1x8e1vnw1vn61x8g1zlu1vn4</Set>
      <Set name="handshakeTimeout">2000</Set>
    </New>
  </Arg>
</Call>
```

3 In the `<Set name="Port">` directive, change the port number.

   When adding the Cloud Manager Orchestration Server as a zone in the Cloud Manager Application Console, you specify this port as the server port.

4 Save the `jetty-ssl.xml` file.

5 Restart the Cloud Manager Web Service:

   ```
   /etc/init.d/novell-pso-ws restart
   ```

8.1.2 Disabling SSL v3 in Java Jetty

Cloud Manager automatically disables the SSL v3 protocol in Java Jetty on the Cloud Manager Application Server, thereby addressing the vulnerability to potential POODLE (Padding Oracle On Downgraded Legacy Encryption) attacks. For more information about POODLE, see *Common Vulnerabilities and Exposures CVE-2014-3566*.

If you need to disable SSL v3 in Java Jetty on your existing Cloud Manager Application Server, you must manually modify the `/opt/netiq/cloudmanager/deploy/jetty/etc/jetty.xml` file:

1 Navigate to the `/opt/netiq/cloudmanager/deploy/jetty/etc/jetty.xml` file and save a copy as `jetty-xml-OLD`.

2 Open the `jetty.xml` file in a text editor.

3 If SSL is enabled, the `jetty.xml` file contains a section that looks like the following. Delete this section from the file.

   NOTE: You must remove the old section. Commenting it out can cause the configuration script to fail when you perform an upgrade.

   ```
   <Call name="addConnector">
     <Arg>
       <New class="org.eclipse.jetty.server.ssl.SslSelectChannelConnector">
         <Set name="Port">[SSL Port Number]</Set>
         <Set name="maxIdleTime">120000</Set>
         <Set name="keystore">
           <SystemProperty name="karaf.home" default="." />[Keystore File Name]
         </Set>
         <Set name="password">[Keystore Password]</Set>
         <Set name="keyPassword">[Key Password]</Set>
         <Set name="wantClientAuth">true</Set>
       </New>
     </Arg>
   </Call>
   ```

4 Replace the old section of the `jetty.xml` file with the following:

```
<Call name="addConnector">
  <Arg>
    <New class="org.eclipse.jetty.server.ssl.SslSelectChannelConnector">
      <Arg>
        <New class="org.eclipse.jetty.http.ssl.SslContextFactory">
          <Set name="keyStore">
            <SystemProperty name="karaf.home" default="." />
            [Keystore File Name]
          </Set>
          <Set name="keyStorePassword">[Keystore Password]</Set>
          <Set name="keyManagerPassword">[Key Password]</Set>
          <Set name="ExcludeProtocols">
            <Array type="java.lang.String">
              <Item>SSLv3</Item>
            </Array>
          </Set>
        </New>
      </Arg>
      <Set name="Port">[SSL Port Number]</Set>
      <Set name="maxIdleTime">120000</Set>
      <Set name="wantClientAuth">true</Set>
    </New>
  </Arg>
</Call>
```

5 Save the changes.

6 Verify that the SSL v3 protocol is disabled.

- The Cloud Manager secure URL (HTTPS) should be functional.
- There should be an entry in the log that shows that the SSLv3 protocol is not in the enabled protocol list. For example:

  ```
  [12 Jan 2015 06:37:08] INFO  | g.ops4j.pax.web) | SslContextFactory   |  96
  | Enabled Protocols [SSLv2Hello, TLSv1, TLSv1.1, TLSv1.2] of [SSLv2Hello,
  SSLv3, TLSv1, TLSv1.1, TLSv1.2]
  ```

8.2 Enabling a Non-Secure Connection

1 On the Cloud Manager Orchestration Server, open the `jetty.xml` file:

  `/etc/opt/novell/pso-ws/jetty/jetty.xml`

2 Uncomment the `<Call name ="addConnector">` section that enables the non-secure port.

3 If you want the Cloud Manager Web Service, which handles the connection for the server, to listen on a port other than 8080, change the port number.

  When adding the Cloud Manager Orchestration Server as a zone in Cloud Manager, you specify this port as the server port.

4 Save the `jetty.xml` file.

5 Restart the Cloud Manager Web Service:

  `/etc/init.d/novell-pso-ws restart`

9 Launching the Orchestration Console and Logging in to the Orchestration Server

When you have installed and configured the Orchestration Server, you can launch and log in to the Orchestration Console.

- Section 9.1, “Launching the Orchestration Console,” on page 71
- Section 9.2, “Logging In Explicitly to a Named Server,” on page 72

NOTE: The Orchestration Console uses TCP port 1099 for its initial connection and then selects a port in the ephemeral port range (ports 32768-65535) for additional communications. If you have problems connecting to the Orchestration Console, ensure that these ports on the server are reachable from the client.

9.1 Launching the Orchestration Console

When the Orchestration Console is launched, it broadcasts throughout the network to discover all of the Orchestration Servers that have been previously installed. The server or servers are displayed at the root of the Explorer panel in the Orchestration Console.

To launch the Orchestration Console:

1. Navigate to the location where the Orchestration Console was installed:
   - **SLES:** Change to the following directory:

     `/opt/novell/zenworks/zos/server/bin`
   - **Windows:** In the **Start** menu, click **All Programs** > **Novell** > **ZOS** > **Clients**.
2. Launch the Orchestration Console:
   - **SLES:** Use the following command to launch the Orchestration Console:

     `./zoc`
   - **Windows:** In the **Start** menu, click **Programs** > **Cloud Manager Orchestration Clients** submenu, then click **Cloud Manager Orchestration Console**.
3. In the Orchestration Console, log in to the Orchestration Server by selecting a server in the Explorer tree.

IMPORTANT: If you are not operating in a broadcast-capable network and you have installed the Orchestration Console on a machine with a different subnet from the server, the console might not be able to discover your Orchestration Server. See “Logging In Explicitly to a Named Server” on page 72 for the login procedure in this scenario.

9.2 Logging In Explicitly to a Named Server

If you do not see the Orchestration Server you want to log into in the Explorer tree, you must log in explicitly to the server you want.

1. In the Orchestration Console, click **Server,** then click **Login** to display the Remote Connection dialog.

   Orchestrate supports multiple servers on the same network.

   This login option allows you to select the server before you enter the administrator name and password.
2. In the dialog, specify the IP address of the Orchestration Server in the **Server Address** field, then click **OK** to display the login dialog.
3. Specify the administrator name (created during the install) in the **Username** field, specify the administrator password in the **Password** field, then click **OK** to log in to the server.

10 Creating a Resource Account

After being installed on a computing node, having its credentials defined, and associating itself with the computing node, the Orchestration Agent begins broadcasting the availability of its host as a potential computing resource. Before the Orchestration Server can allow an agent to authenticate and establish ongoing communication, you need to create a resource account for the agent on the Orchestration Server. When this account is created or “registered,” the agent’s host node can be discovered and recognized as a computing resource that can perform the jobs assigned to it.

It is also possible to create a resource account for an agent before that agent is actually installed on a computing node.

You can create a resource account on the Orchestration Server and have it waiting in an offline state in anticipation of agent installation and login.

This section includes the information you need to create a resource account on the Orchestration Server:

- Section 10.1, “Opening the Resources Monitor,” on page 74
- Section 10.2, “Automatically Registering a Resource,” on page 74
- Section 10.3, “Selecting a Resource for Manual Registration,” on page 75
- Section 10.4, “Manually Registering a Resource in the Orchestration Console,” on page 76

When resources are created, connected to the Orchestration Server and online, a provisioning adapter job deploys and runs a discovery process on its own.

For information about manually configuring a resource account, see Appendix 10.4, “Manually Registering a Resource in the Orchestration Console,” on page 76.

10.1 Opening the Resources Monitor

Now that you have installed an Orchestration Server and launched the Orchestration Console, you can begin to create resource accounts.

1 Open the Orchestration Console and click **Resources** to open the Resources Monitor in the admin view of the Orchestration Console.

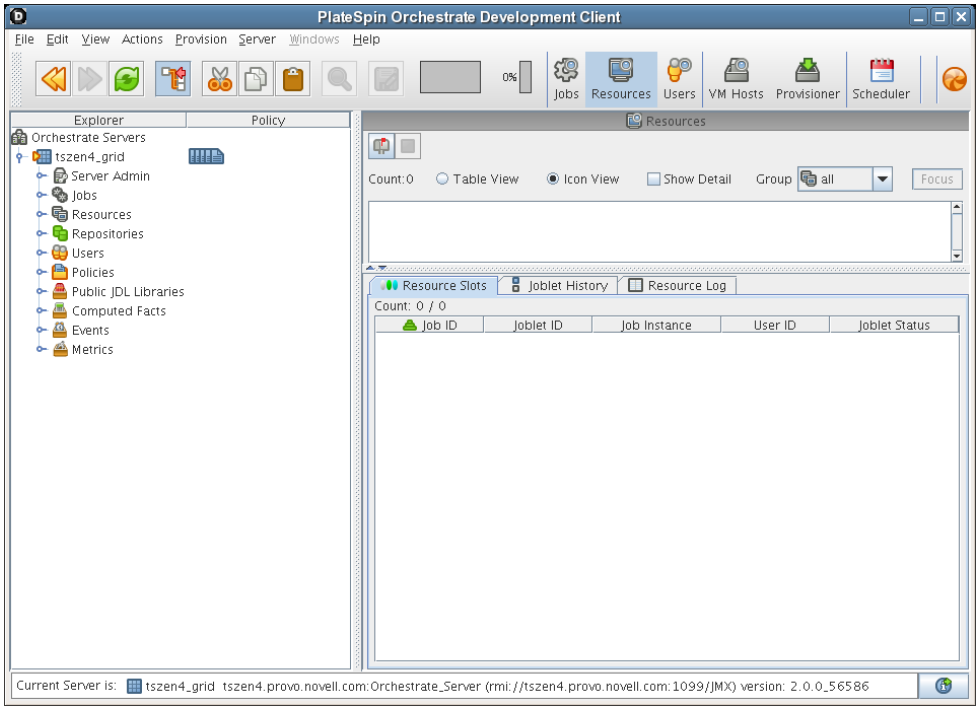

From this monitor, you can see the resources that are connected to the server and what they are doing in the grid.

If an agent is installed but has not been registered (that is, no account is created for it), it attempts a server login every 90 seconds. If this is the case (as in the figure above), the Resource Registration icon has a “flag up”  status, meaning that an agent is waiting to register. If the icon has a “flag down”  status, either no Orchestration Agents have been installed in the network or all active agents are logged in, so none are waiting to register.

The Resources Monitor has many features to help you manage resources when they are registered, including the jobs and joblets assigned to individual resources.

10.2 Automatically Registering a Resource

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 resource account, you can do so at the Orchestration Console.

1 In the Orchestration Console, select the grid object in the Explorer tree to open the **Authentication** page in the admin view.

2 In the **Resources** section of the page, select the **Auto Register Agents** check box, then click the Save icon  in the toolbar to save the setting.

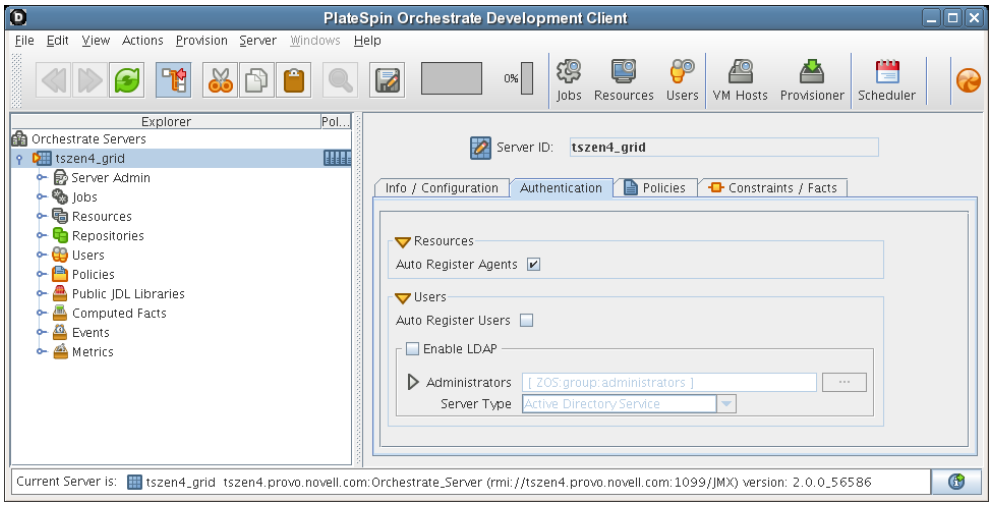

The resource object is created and registered in the Orchestration Server, although it is offline (the object is dimmed in the tree of the Explorer panel) until it the agent tries to log in.

The next time the agent tries to log in, it is automatically authenticated and the Orchestration Server creates a a new resource account.

When the resource is online, the Resources Monitor displays a labeled box representing the registered agent. This box includes information about the agent, including the number of available slots it has and a status color indicating its state of readiness for Cloud Manager jobs.

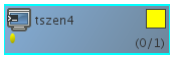

The status color window can be white (inactive), green (available for use), or blue (in use). If the color changes from green to blue, a job is running on this resource. To find out what kind of job is running, you can click the **Jobs** monitor button on the toolbar.

10.3 Selecting a Resource for Manual Registration

If you do not select the **Auto Register Agents** check box on the grid object's **Authentication** page, you have the option of explicitly accepting or denying the login attempts of a resource, thus preventing it from creating an account.

The following steps assume that you have already created a resource in your grid.

1. In the Resources Monitor, click the Resource Registration (mailbox) icon to open the Resource Registration Monitor dialog.

   This dialog lets you preview the Orchestration Agents that are installed in the network and trying to log in to the server. The top row of radio buttons is a mass selector for all listed agents, allowing you the choice to accept, deny, or ignore automatic registration for all agents, both those currently listed and those that might try to log in later.

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

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

3 From the Orchestration Console, open the Resources Monitor to observe the resource object you created change from offline to online. When the object is no longer dimmed, the agent has logged in as a resource and is registered.

When the resource is online, the Resources Monitor displays a labeled box representing the registered agent. This box includes information about the agent, including the number of available slots it has and a status color indicating its state of readiness for Orchestrate jobs.

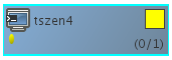

The status color window can be white (inactive), blue (in use), green (available for use), or blue (in use). If the color changes from green to blue, a job is running on this resource. To find out what kind of job is running, you can click on the **Jobs** monitor button on the toolbar.

10.4 Manually Registering a Resource in the Orchestration Console

If you want a higher level of security between the agent and the server, you can manually create a resource account in the Orchestration Console before the Orchestration Agent is installed. This section walks through both stages of the procedure.

- Section 10.4.1, “Using the Orchestration Console to Create a Resource Account,” on page 76
- Section 10.4.2, “Installing an Orchestration Agent to Match the New Resource,” on page 77

10.4.1 Using the Orchestration Console to Create a Resource Account

Use the following steps to create a resource object in the Orchestration Console.

1 Ensure that the **Auto Register Agents** check box on the grid object’s **Authentication** page is not selected (see Step 2 on page 74).

2 (Optional) Create a new resource from the Explorer panel in the Orchestration Console:

2a In the Explorer panel in the Orchestration Console, right-click **Resources**, then click **New Resource** to display the Create a new Resource dialog.

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

3 (Optional) Create a new resource from the Main Menu in the Orchestration Console.

3a In the Orchestration Console, click **Actions** > click **Create Resource** to display the an expanded version of the Create a new Resource dialog.

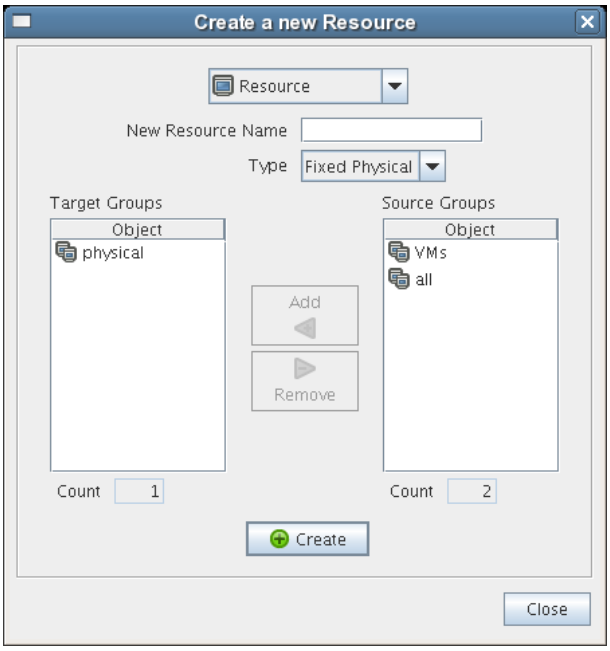

This dialog includes a method for designating the resource as a fixed physical type or a virtual machine type. It also includes a method for including the resource in various resource groups. In this walkthrough, we will install an Orchestration Agent on a fixed physical resource and include it in the **physical** resource group.

The Virtual Machine resource type is not available if you installed the High Performance Computing license only for Orchestrate.

3b Ensure that **Fixed Physical** is selected in the **Type** drop-down box, specify the new resource name in the **New Resource Name** field, click **Create**, then click **Close**.

The resource account is created, but is offline , as indicated by its object icon in the Explorer panel or in the Information view of each resource group to which it belongs. The resource is not online until an Orchestration Agent matching the resource is installed.

10.4.2 Installing an Orchestration Agent to Match the New Resource

This section demonstrates installing an Orchestration Agent to be used as a resource in your Orchestration grid. The information in this part of the walkthrough assumes that a resource account has already been created for the Orchestration Agent being installed.

1. From the managed device desktop, launch a browser to access the web page for Orchestrate, as described in Section 5.2.1, “Obtaining the Agent Installer and Supporting Files from the Administrator Information Page,” on page 42.
2. Scroll to the **Installation** section of the page:

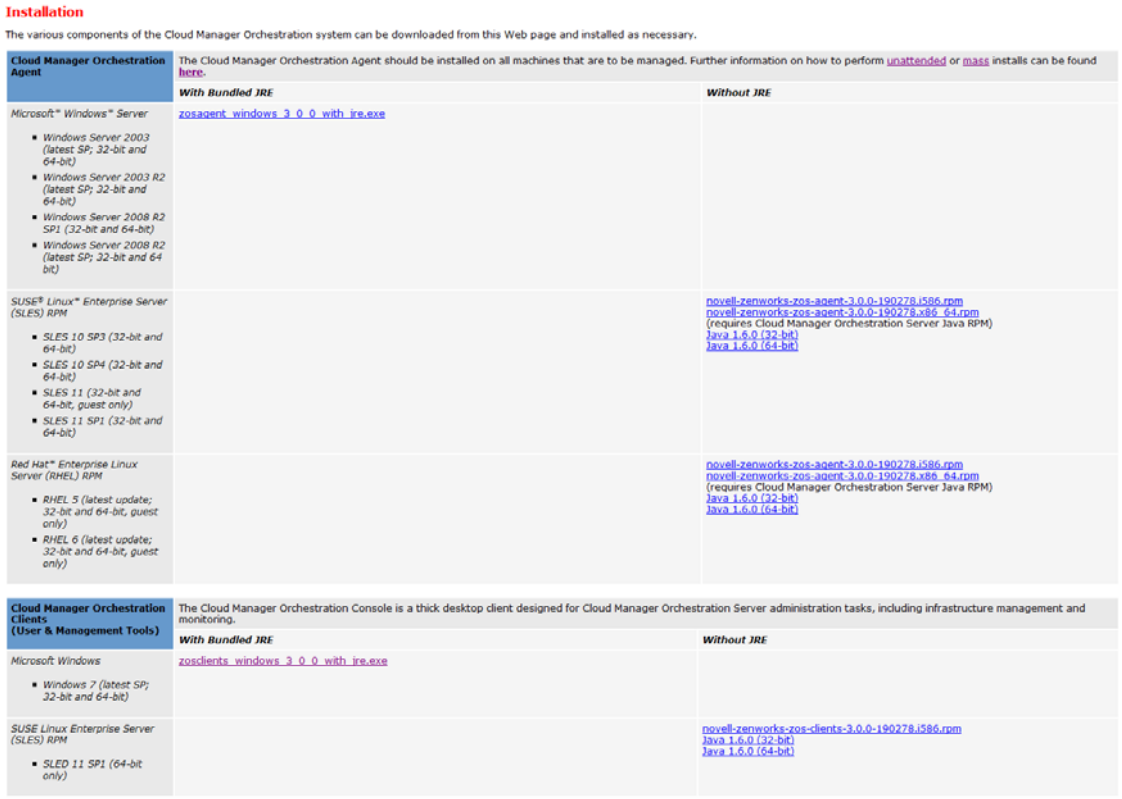

Installation

The various components of the Cloud Manager Orchestration system can be downloaded from this Web page and installed as necessary.

Cloud Manager Orchestration Agent	The Cloud Manager Orchestration Agent should be installed on all machines that are to be managed. Further information on how to perform unattended or mass installs can be found here.	
	With Bundled JRE	***Without JRE***
Microsoft® Windows® Server ▪ *Windows Server 2003 (latest SP; 32-bit and 64-bit)* ▪ *Windows Server 2003 R2 (latest SP; 32-bit and 64-bit)* ▪ *Windows Server 2008 R2 SP1 (32-bit and 64-bit)* ▪ *Windows Server 2008 R2 (latest SP; 32-bit and 64 bit)*	zosagent_windows_3_0_0_with_jre.exe	
SUSE® Linux™ Enterprise Server (SLES) RPM ▪ *SLES 10 SP3 (32-bit and 64-bit)* ▪ *SLES 10 SP4 (32-bit and 64-bit)* ▪ *SLES 11 (32-bit and 64-bit, guest only)* ▪ *SLES 11 SP1 (32-bit and 64-bit)*		novell-zenworks-zos-agent-3.0.0-190278.i586.rpm novell-zenworks-zos-agent-3.0.0-190278.x86_64.rpm (requires Cloud Manager Orchestration Server Java RPM) Java 1.6.0 (32-bit) Java 1.6.0 (64-bit)
Red Hat™ Enterprise Linux Server (RHEL) RPM ▪ *RHEL 5 (latest update; 32-bit and 64-bit, guest only)* ▪ *RHEL 6 (latest update; 32-bit and 64-bit, guest only)*		novell-zenworks-zos-agent-3.0.0-190278.i586.rpm novell-zenworks-zos-agent-3.0.0-190278.x86_64.rpm (requires Cloud Manager Orchestration Server Java RPM) Java 1.6.0 (32-bit) Java 1.6.0 (64-bit)

Cloud Manager Orchestration Clients (User & Management Tools)	The Cloud Manager Orchestration Console is a thick desktop client designed for Cloud Manager Orchestration Server administration tasks, including infrastructure management and monitoring.	
	With Bundled JRE	***Without JRE***
Microsoft Windows ▪ *Windows 7 (latest SP; 32-bit and 64-bit)*	zosclients_windows_3_0_0_with_jre.exe	
SUSE Linux Enterprise Server (SLES) RPM ▪ *SLED 11 SP1 (64-bit only)*		novell-zenworks-zos-clients-3.0.0-190278.i586.rpm Java 1.6.0 (32-bit) Java 1.6.0 (64-bit)

3 In the agent section of the Administrator Information page, find the installer link for the operating system of the device where you want to install the agent. For this walkthrough, we will install the Windows agent on a Windows operating system.

4 Click the installer link to download the `zosagent_windows_<version>_with_jre.exe` version of the agent to the computing node where you plan to install it.

5 From the machine where you will install the agent (in this walkthrough, a Windows Server 2008 64-bit machine), open the desktop and navigate to the location where you saved the Orchestration Agent file, then double-click the `zosagent_windows_<version>_with_jre.exe` icon to launch the Orchestration Agent Setup Wizard.

6 Follow the prompts in the wizard until the **Identify Orchestration Server** page displays, then ensure that you correctly enter the *Cloud_Manager_Orchestration_Server_name* in the **Orchestration Server** field.

IMPORTANT: Ensure that the name you give the agent during the installation matches the name of the resource account you created in "Using the Orchestration Console to Create a Resource Account" on page 76.

You might find it easier to click **Discover** so that the installer searches for and finds the Orchestration Server on the network.

7 Accept the remaining defaults on the wizard pages to complete the installation of the Agent.

8 When the installation is complete, click **Finish** to exit the wizard.

9 In the Orchestration Console, open the Resources Monitor to observe the resource object you created change from offline to online. When the object is no longer dimmed, the agent has logged in as a resource and is registered.

When the resource is online, the Resources Monitor displays a labeled box representing the registered agent. This box includes information about the agent, including the number of available slots it has and a status color indicating its state of readiness for Orchestrate jobs.

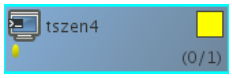

The status color window can be white (inactive), green (available for use), or blue (in use). If the color changes from green to blue, a job is running on this resource. To find out what kind of job is running, you can click on the **Jobs** monitor button on the toolbar.

11 Configuring the vSphere Provisioning Adapter

After you install the Cloud Manager Orchestration Server and configure its components, you are ready to preparing it to receive information about the VMs it discovers in your vSphere Server environments. The Cloud Manager vSphere Provisioning Adapter (`vsphere`) detects changes that occur in the vSphere Server environment and synchronizes many of them to the Cloud Manager Orchestration Server, which in turn communicates them to the Cloud Manager Application Server and ultimately to the Web Console.

For additional information about the vSphere provisioning adapter, see "The VMware vSphere Provisioning Adapter" in the *Cloud Manager Administrator Reference*.

- Section 11.1, "Initializing the vSphere Provisioning Adapter on the Orchestration Server," on page 81
- Section 11.2, "Managing the vSphere Provisioning Adapter," on page 83
- Section 11.3, "Configuring the vSphere Provisioning Adapter for VMware DRS Clusters and MetroClusters," on page 87
- Section 11.4, "Discovering Enterprise Resources in Multiple vSphere Environments," on page 90
- Section 11.5, "Discovering vSphere Hosts and VMs," on page 93

11.1 Initializing the vSphere Provisioning Adapter on the Orchestration Server

You must configure a `vsphere` provisioning adapter job that authenticates to a VMware vSphere environment. After you configure the adapter, you can discover VMs in that environment.

To configure multiple vSphere environments, see Section 11.4, "Discovering Enterprise Resources in Multiple vSphere Environments," on page 90.

1. Ensure that the Orchestration Agent is installed and started on a supported VMware vSphere host. See Chapter 5, "Installing Cloud Manager Orchestration Components," on page 39.

   You can install the agent either locally on the vCenter Server (the vCenter appliance is not supported), or on a dedicated system (virtual or physical) as long as the OS in that system is supported for the Orchestration Agent.

   NOTE: If you want to connect multiple vCenter Server environments, see Section 11.4, "Discovering Enterprise Resources in Multiple vSphere Environments," on page 90.

2. In the Cloud Manager Orchestration Console, log in to the Orchestration Server that you want to use to manage vSphere VMs.
3. In the Explorer tree of the Orchestration Console, select the Orchestration Server or "grid" object, then select the **Authentication** tab in the admin view to open the Authentication page.

4 Create a credential to authenticate to the vCenter Server. In most cases, the credential is for "administrator" account of the Windows machine where the vCenter Server is running.

4a On the Authentication page, scroll to the Credential Manager (consisting of the **Stored Credentials** panel and the **Stored Certificates** panel), then click **Add Credential** to display the **Edit Credential** dialog.

4b Fill in the fields of the dialog:

- **Name:** Enter a value in that identifies the vCenter Web Service to log in to.
- **User:** Enter the user name used to connect to the vCenter Web Service.
- **Secret:** Enter the password of the vCenter Web Service user.
- **Type:** (Optional) Enter any string that lets you categorize similar credentials into a category or group. For example, for the vSphere provisioning adapter you might enter a "type" called `vsphere`.

For more information, see "Authentication Page" in the.*Cloud Manager Administrator Reference*.

4c Click **Add**.

5 In the Explorer tree, expand the **Jobs** container, then expand the **provisionAdapters** container to expose all of the provisioning adapter jobs.

6 In the Explorer tree, select the **vsphere** job to open the Admin view.

7 In the Admin view, select the **Job Configuration** tab to open the Job Configuration page, then expand the Accounts table on this page.

8 On the Accounts table, select **Add** to open the Add a New Account dialog.

8a Fill in the fields of the dialog:

- **Account Name:** Enter the name you wish to use to refer to this vCenter server as within Orchestrator. This name can be any value, but once selected, it should not be changed.
- **vSphere Web Service URL:** Enter the URL of the vCenter Web Service server.

  Syntax: `https://address-of-vcenter-server/sdk`

  Example: `https://vcenter.server.test/sdk`, where `vcenter.server.test` is the fully qualified domain name (FQDN) of the vCenter server. You can use the IP address rather than the FQDN.
- **Credential Name:** Enter the name of the credential from the Credential Manager that you want to use for logging in to the vCenter Web Service server.
- **Auto Portgroup Creation:** (Optional) If it is selected and the `vsphere_ignoreNetwork` policy is used, port groups are automatically created on a host if it does not have access to the specified network.
- **Auto Portgroup Disconnect:** (Optional) If it is selected, the vNIC on a VM is disconnected when it is shut down.
- **Auto Portgroup Deletion:** (Optional) If it is selected, when the VM is shut down, it checks for port groups on the VM host that has no VMs associated with it and deletes them, if possible. This setting is best used with **Auto Portgroup Creation** and **Auto Portgroup Disconnect**.

9 Associate the `vsphere_client` policy to the resource that will access the vCenter Server.

When the vsphere provisioning adapter job starts, this policy constrains the resource for the job to run the Web Service commands.

9a In the Explorer tree, expand the **Resources** group, select the client resource that is to access the vCenter Web Service server, then select the **Policies** tab in the admin view to open the **Policies** page.

9b On the Policies page, click **Choose** to open the Policy Selection dialog, then in the **Source Policies** list, select the **vsphere_client** policy, click **Add**, then click **OK** to associate this policy with this resource.

For more information, see "Resource Policies Page" in the *Cloud Manager Administrator Reference*.

After you have completed the authentication configuration, continue with "Running Discovery for a vSphere Environment" on page 93.

11.2 Managing the vSphere Provisioning Adapter

After you set up the vSphere Provisioning Adapter, you can modify the configuration to manage VMs in your environment.

- Section 11.2.1, "VM Management Policies for the vSphere Provisioning Adapter," on page 83
- Section 11.2.2, "Synchronizing vSphere vCenter System Changes to the Orchestration Server," on page 84
- Section 11.2.3, "Assigning a vSphere VM to a Resource Pool," on page 84
- Section 11.2.4, "Setting Up Orchestration VNC for a VM Managed by vSphere," on page 85

11.2.1 VM Management Policies for the vSphere Provisioning Adapter

The following table provides detailed information about policies associated with the vSphere provisioning adapter that are used to manage the vSphere hosts and the VMs in the grid. The policy settings are applied to all the VMware VMs in the grid.

Table 11-1 *Virtual Machine Management Policies for vSphere*

Policy Name	Description	Additional Details
vsphere	Contains the constraints used to select the vCenter Server resources.	Do not modify this policy.
vsphere_assignPool	If you need to assign the VMs to a certain cluster (for example, a cluster root pool), or if you want to assign VMs to pools "owned" by your customers, use this policy.	When applied this policy allows the VM to reside only on VM hosts that have access to the assigned resource pool (`resource.vm.pool`).
vsphere_client	Contains the settings used to run the vsphere job on the associated vSphere resource.	You need to associate the vsphere_client policy to a vSphere resource before the discovery works. For more information, see Step 9 on page 82.

Policy Name	Description	Additional Details
vsphere_ignoreNetwork	Includes special facts that allow VMs to consider a VM host despite a missing required Network.	If ignoreNetworkCheck is set, a vBridge (portgroup) can be dynamically created on a VM power-on event. This works in conjunction with the auto_portgroups_creation fact found in the vpshere.policy. Ensure that you set the `auto_portgroups_creation` fact to true or else the portgroup will not be created during the VM power-on event.
vspherePA	Includes the basic constraints for the vsphere provisioning adapter.	Do not modify this policy.
vSphereUpdate	Includes settings for the vsphereUpdateDaemon job. The policy can be modified directly or the user can edit job arguments in the schedule that is created by default installation of the Cloud Manager Orchestration Server.	For more information, see “Configuring the vSphereUpdate Client” on page 97.
vsphereVmHostVnc	Includes port settings to identify a range of ports to be used for remote connections on a specified VM host.	When applied, this policy defines a range of port numbers to be used for remote connections. As VMs are provisioned, they are assigned a port number within the configured range for remote access. This applies only when the VNC mode is **automatic** (the default) as defined in the vsphereVnc policy.
vsphereVnc	Includes a setting to allow remote desktop connections to vSphere VMs.	For more information, see “Setting Up Orchestration VNC for a VM Managed by vSphere” on page 85.

11.2.2 Synchronizing vSphere vCenter System Changes to the Orchestration Server

The vSphereUpdate tool connects to the vSphere vCenter system to get the changes made outside of Cloud Manager, and synchronize them to the appropriate objects in Cloud Manager. After you congure the provisioning adapter and discover the vSphere Server environments, set up vSphereUpdate jobs for them. See Chapter 12, “Configuring vSphereUpdate Client and Monitoring Jobs,” on page 95.

11.2.3 Assigning a vSphere VM to a Resource Pool

All VMs managed by vSphere are assigned to either the default (named “Resources”) or a named resource pool. When vSphere VM images are discovered by the Orchestration Agent, the `resource.vm.pool` fact for each VM is set with what is known by vSphere as a “pool assignment.”

If you do not need to restrict VMs based on resource pool assignment, then no policy configuration is necessary and you can provision the VMs as usual, but if you need to assign the VMs to a certain cluster (for example, the cluster root pool), or if you want to assign VMs to pools “owned” by your customers, you can configure the `vsphere_assignPool` policy to accomplish this.

To ensure that the Orchestration Server always provisions a VM to the resource pool where that VM resides:

1. Assign the `vsphere_assignPool` policy to the VM or a group of VMs. No changes to the actual policy file are necessary.

   During provisioning of the VM, the Orchestration Server verifies and relocates the VM (as necessary) to maintain the validity of the pool assignment.

2. (Conditional) If the VM does not reside in the correct resource pool, look up the ID of the resource pool in vCenter and modify the `resource.vm.pool` fact to reflect the correct pool assignment. The Orchestration Server relocates the VM to the specified resource pool at the next provision.

   Alternatively, use vSphere to move the resource to the proper pool and re-run the VM discovery process.

NOTE: Because Cloud Manager deploys workloads managed by the VMware ESX hypervisor only into resource pools that have been configured on an associated cluster, the Cloud Manager administrator must assign these VMware resource pools to Cloud Manager resource groups. For more information, see "Creating Resource Groups" in the *Cloud Manager Procedures Guide*.

11.2.4 Setting Up Orchestration VNC for a VM Managed by vSphere

When you right-click a VM resource, you have the option of launching a remote virtual network computing (VNC) session console of that VM's desktop. This section provides information about setting up the Orchestration Server to accommodate a VNC session for a VM.

ESX 4.*x* servers managed by vSphere might have a firewall in place to protect some ports from being open or closed. The vsphere provisioning adapter opens the appropriate ports to accommodate VNC connections from a remote console. These ports are opened when the Orchestration Server discovers the servers. This is not true for ESX 5.*x* servers managed by vSphere, where the ports require manual opening.

NOTE: With the vSphere 5 release, VMware removed **VNC Server** as a service than can be directly administered by using the VMware Client or the VMware Client libraries and APIs. Although the VNC functionality still works on ESXi servers, the firewall must be opened to allow access.

For information about enabling VNC access for ESXi 5 servers, see "Enabling VNC Access to vSphere 5 VM Guest Consoles" in the *Cloud Manager Administrator Reference*.

Although you can change these settings at any time, they take effect for a vSphere VM after a non-running VM is provisioned, or after you perform an **Apply Config** action on a running VM.

To set up VNC session connectivity for the VM managed by the `vsphere` provisioning adapter job:

1. In the Explorer tree of the Orchestration Console, select the Grid Server object where you are logged in, then select the **Authentication** tab to open the server's authentication page.
2. In the **Stored Credentials** panel (also known as the "Credential Manager") of the Authentication page, click **Add Credential** to open the Add Credential dialog.
3. In the Add Credential dialog, specify a credential that includes the VNC password you want to use, then click **Add**. List the credential type as `vnc`.

Although a user is required when you create the credential, this value is not used in the remote session. Only the **secret** field is used when making the connection.

4 Configure the vsphereVNC policy.

4a In the Explorer tree, expand the Policies folder to display the list of policies, then select **vsphereVNC** policy to open the Policy Editor.

4b In the Policy Editor, modify the `vnc.credential` fact value to be the name of the credential you created in Step 3, then click the **Save** icon.

Modifying this policy is not necessary unless you want to assign the same credential to every vSphere VM or groups of vSphere VMs. Otherwise, you can select the credential on a per-VM basis from the **VNC Credential** drop-down list on the Resource Information panel of the VM's Info/Groups page.

5 (Conditional) Configure the vsphereVmHostVnc policy for a VM host.

Modifying this policy is not necessary unless the `resource.vnc.mode` fact of the vsphereVNC policy is set to **automatic**. When the ports defined in this range have been consumed, further vSphere VM provisioning fails.

The default port range in the policy is 1000-9999. If you want to provide remote capabilities to more VMs on a host or cluster, you need to alter the policy configuration to add more ports to the range. You can also reconfigure the policy to use a different range of ports.

5a In the Explorer tree, expand the Policies folder to display the list of policies, then select **vsphereVmHostVnc** policy to open the Policy Editor.

5b In the Policy Editor, modify the `vpshere.port.min` fact value as the lower end of the range of ports you want to be used as remote connections for this VM host.

5c In the Policy Editor, modify the `vpshere.port.max` fact value as the upper end of the range of ports you want to be used as remote connections for this VM host, then click the **Save** icon.

6 Associate the vsphereVNC policy to a VM Resource Group or VM.

6a In the Explorer tree, select the VM Resource Group (or an individual VM) managed by the vsphere provisioning adapter, then in the admin view, select the **Policies** tab to open the Policies page for this group.

6b On the Policies page, select **Choose** to open the Policy Selection dialog.

6c In the **Source Policies** list of the Policy Selection dialog, select the **vsphereVnc** policy, click **Add** to move it to the associated **Policies** list, then Click **OK**.

7 Associate the vsphereVmHostVnc policy to a VM host.

7a In the Explorer tree, select the VM host managed by the vsphere provisioning adapter, then in the admin view, select the **Policies** tab to open the Policies page for this group.

7b On the Policies page, select **Choose** to open the Policy Selection dialog.

7c In the **Source Policies** list of the Policy Selection dialog, select the **vsphereVMHostVnc** policy, click **Add** to move it to the associated **Policies** list, then Click **OK**.

8 On the Orchestration Console Menu Bar, click **Provision** > **Discover VM Hosts and Repositories**.

In vSphere 4.*x* environments, this action opens or closes the firewall on the VM hosts to allow VNC access. This access is based on the `vsphere.openVncFirewallPort` fact in the vsphere policy.

For ESX 5.*x* servers managed by vSphere, the ports require manual opening. See "Enabling VNC Access to vSphere 5 VM Guest Consoles" in the *Cloud Manager Administrator Reference*.

9 (Conditional: For VMs that are running) From the Explorer tree, right-click a vSphere-managed VM, then select **Apply Config**.

If the VM for which you want to open a VNC session is not running, simply reprovision the VM.

10 If the vSphereUpdate Client is running for your vCenter server, refresh the Orchestration Console.

or

If the vSphereUpdate Client is not running for your vCenter server, right-click the VM object and select **Resync State**.

If you don't want to resync before using the VNC console, ensure that you configure the vSphereUpdate Client beforehand. See "Configuring the vSphereUpdate Client" on page 97.

11 Right-click the VM object and select **Launch Remote Desktop** to open the login dialog for the VNC session.

12 In the login dialog, enter the VNC password that you created in the Credential Manager in Step 3.

The following table lists the VNC-related facts in the vsphere provisioning adapter and provides a description of each of those facts.

Table 11-2 *vSphere VNC Facts*

Fact Name	Description
`resource.vnc.ip`	The IP address of the VM host where the VM is running
`resource.vnc.port`	The port currently assigned to the VM. The value is -1 if VNC is disabled for the VM.
`resource.vnc.credential`	The credential containing the VNC password. This is the name of the credential itself, not the user name or the password contained in the credential.
`resource.vnc.mode`	Determines how VNC port assignments are handled. This value must be `automatic`, `manual`, or `off`. ◆ If mode = `automatic`: the Orchestration Server attempts to select the next available VNC port. ◆ If mode = `manual`: The port value specified in the VM's `resource.vnc.port` fact is used. ◆ If mode = `off`: The VNC console is disabled.
`resource.remotedesktop`	Controls enabling or disabling the Launch Remote Desktop action in the Orchestration Console.

11.3 Configuring the vSphere Provisioning Adapter for VMware DRS Clusters and MetroClusters

After you set up the vSphere Provisioning Adapter, you can modify the configuration to discover VMware DRS clusters in your environment.

- Section 11.3.1, "Setting Up Orchestration to Accommodate VMware DRS Clustering and Updates," on page 88
- Section 11.3.2, "Constraining vSphere VMs to Their Assigned Resource Pools," on page 88

- Section 11.3.3, “Configuring the Orchestration Server to Limit Datastore Visibility in vSphere Clusters,” on page 88
- Section 11.3.4, “Configuring Cloud Manager Orchestration for a vSphere MetroCluster Environment,” on page 89

11.3.1 Setting Up Orchestration to Accommodate VMware DRS Clustering and Updates

The Orchestration Server supports the discovery of VMware vSphere clusters used for high availability in a VMware environment or managed by the VMware Distributed Resource Scheduler (DRS) after an Orchestration Agent has been deployed into such an environment. In this scenario, Cloud Manager Orchestration also lets you verify when actions have taken place outside of Cloud Manager, such as when DRS moves a VM to an alternate host in the cluster or when an administrator moves a VM into a different resource pool.

Any vSphere clusters discovered by Cloud Manager and managed by DRS are listed in the Orchestration Console as members of a convenience group (for example, a group named `clusters_`*`vsphere`*).

You can learn about read-only cluster-related facts for the following aspects of discovered clusters in “Unique VM Host Cluster Facts” in the *Cloud Manager Administrator Reference*:

- VM Host Cluster object
- VM Host residing in a cluster
- VMs hosted in a cluster

11.3.2 Constraining vSphere VMs to Their Assigned Resource Pools

To assign the VMs to a certain cluster (for example, the cluster root pool), or if you want to assign VMs to a pool “owned” by your customers, configure the vsphere_assignPool policy to a VM or a group of VMs.

1. In the Orchestration Console tree view, select the VM or Group of VMs that you wish to constrain to their assigned resource pool.
2. In the admin view, select **Policies** to open the Policies page.
3. On the Policies page, select **Choose** to display the Policy Selection dialog.
4. In the Source Policies list, select **vsphere_assignPool**, click **Add** to move it to the Associated Polices list, then click **OK**.

11.3.3 Configuring the Orchestration Server to Limit Datastore Visibility in vSphere Clusters

If you want to limit the number of datastores (that is, repositories that are modeled in the Orchestration Server) that are available to a vSphere cluster, you can assign a policy similar the policy below to the undesired repository or repositories:

```
<policy>
  <repository>
    <fact name="enabled" type="Boolean" value="False" />
    <fact name="provisioner.jobs">
      <array type="String">
      </array>
    </fact>
  </repository>
</policy>
```

This disables the repository for use with the cluster.

11.3.4 Configuring Cloud Manager Orchestration for a vSphere MetroCluster Environment

Your VMware cluster environment might be configured to accommodate a MetroCluster (sometimes called a *stretch cluster*) high availability solution. MetroCluster allows for synchronous mirroring of server volumes between two separate storage controllers, usually in the same datacenter. When a MetroCluster is in place, vCenter sees two available repositories, even though the location on the ESX server (that is, the VM host) is unknown.

By using a special job called `metroCluster`, the Cloud Manager Orchestration Server associates all workloads provisioned to a specified repository with its corresponding DRS cluster rule. The Orchestration Server uses this job to work with these rules and facts to determine the DRS group that the VM should be added to. The job ensures that the provision plan balances the VM load fairly across the two repositories.

To implement the MetroCluster compatibility feature in Cloud Manager Orchestration:

1. In the Explorer tree navigate to **Jobs** > **examples**, then right-click **examples** and select **Deploy Job** to open the job deployment dialog, then select the metroCluster job and click **OK**.

   or

   Use the zosadmin command line interface to invoke the following commands for deploying the metroCluster job:

   `zosadmin login <server_hostname_or_server_IP>`

   `zosadmin deploy /opt/novell/zenworks/zos/server/examples/metroCluster.job`

2. From the **Provision** menu, run the **Discover VM Hosts & Repositories** action to discover any new cluster rules or to update existing cluster rules.

   A DRS cluster rule (defined by the vSphere client) would have been created to ensure that VMs in a specified group (for example, `site_a_vms`) run on hosts in a specified group (for example, `site_a_hosts`).

3. Create a new policy to define the value for the new fact.

   The fact should contain the name of the desired cluster rule as its value.

   3a In the Explorer tree, click **Actions** > **Create** > **Create Policy** to open the Create a New Policy dialog.

   3b In the dialog, provide a name for the policy, then click **Create** > **Close** to create the new policy template in the Policy Editor.

   3c In the Policy Editor, substitute the following:

```
<policy>
  <repository>
    <fact name="vsphere.cluster.metro.rule"
          type="String"
         description="The cluster rule to use for VMs associated with this
repository."
          value="foo" />
  </repository>
</policy>
```

NOTE: In this example, you need to replace the `foo` value with the actual name of the DRS cluster rule.

4 Associate the policy with the primary cluster repository.

4a In the Explorer tree, locate and select the primary cluster repository object, then select **Policies** to open the policies page.

4b On the Policies page, select **Choose** to open the Policy Selection dialog.

4c From the Source Policies list, select the name of the policy you created in Step 3 click **Add** to add it to the associated policies list, then click **OK**.

5 Repeat Step 4 on the mirrored repository of the MetroCluster.

When a VM is provisioned (that is, started) or moved to a repository, it should be joined to the cluster rule, which should allow the failover.

11.4 Discovering Enterprise Resources in Multiple vSphere Environments

A data center administrator running VMware products might organize the virtual resources in his or her enterprise into several different vSphere environments. The Cloud Manager Orchestration Server lets you discover and manage all of these enterprise VMs, discovering each relevant VM host, network, repository, and VM within the several vSphere environments and modeling them as objects in the Orchestration Console.

- Section 11.4.1, “Creating Accounts for Each vCenter Environment,” on page 90
- Section 11.4.2, “Configuring the vsphere.vcenters Fact to Include All Accounts Representing a vCenter Server,” on page 91
- Section 11.4.3, “Optionally Specifying an Authentication Certificate for Each vCenter Server,” on page 92

11.4.1 Creating Accounts for Each vCenter Environment

1 In the Explorer tree, select the **vsphere** provisioning adapter job to open the Admin view of this job.

2 Select the **Job Configuration** tab to open the Job Configuration page, then expand the Accounts table on this page.

3 On the Accounts table, select **Add** to open the Add a New Account dialog.

3a Fill in the fields of the dialog:

- **Account Name:** This should match the name of the VCenter environment you are connecting to.
- **vSphere Webservice URL:** Enter the URL of the vCenter Web Service server.

- **Credential Name:** Enter the name of the credential from the Credential Manager that you want to use for logging in to the vCenter Web Service server.
- **Auto Portgroup Creation:** (Optional) If selected and the `vsphere_ignoreNetwork` policy is used, port groups are automatically created on a host if it does not have access to the specified network.
- **Auto Portgroup Disconnect:** (Optional) If selected, the vNIC on a VM is disconnected when it is shut down.
- **Auto Portgroup Deletion:** (Optional) If selected, when the VM is shut down, it checks for port groups on the VM host that has no VMs associated with it and deletes them, if possible. This setting is best used with **Auto Portgroup Creation** and **Auto Portgroup Disconnect**.

4 Repeat Step 3 for every vCenter Server you want to connect to for VM discovery in that vSphere environment.

After you have created Orchestration accounts for each of the vCenter servers in your enterprise, you can continue with “Configuring the vsphere.vcenters Fact to Include All Accounts Representing a vCenter Server” on page 91.

11.4.2 Configuring the vsphere.vcenters Fact to Include All Accounts Representing a vCenter Server

The `vsphere.vcenters` fact can be set to include the definition for all the vCenter server accounts that you identified in “Creating Accounts for Each vCenter Environment” on page 90. This is required to ensure that only certain agents communicate with certain vSphere accounts. You can set this fact in the vsphere_client policy of the vsphere provisioning adapter or by using a policy to apply to the individual Orchestration Agents you installed in your respective vSphere environments.

- “Configuring the vsphere.vcenters Fact in the vsphere_client Policy” on page 91
- “Creating a Policy to Apply to Each Orchestration Resource in the Respective vSphere Environments” on page 92

After you have used one of these methods, continue with “Optionally Specifying an Authentication Certificate for Each vCenter Server” on page 92.

Configuring the vsphere.vcenters Fact in the vsphere_client Policy

You can associate the `vsphere.vcenters` fact in the vsphere_client policy to the resources that access the respective vCenter Servers. When the vsphere provisioning adapter job starts, the policy applies the `vsphere.vcenters` fact to constrain the identified resources for the job to run the Web Service commands.

Use these steps to configure the `vsphere.vcenters` fact:

1 In the Explorer tree, expand the **Policies** group, then select the **vsphere_client** policy to open the **Policy Editor** page in the admin view.

2 In the Policy Editor, scroll to or search for the `vsphere.vcenters` fact, then uncomment it and enter a string value in the array, using the **Account Name** for each vCenter Server (identified in “Creating Accounts for Each vCenter Environment” on page 90) as a string value.

3 In the Explorer tree, expand the **Resources** group, select the client resource that is to access the vCenter Web Service server, then select the **Policies** tab in the admin view to open the **Policies** page.

4 On the Policies page, click **Choose** to open the Policy Selection dialog, then in the **Source Policies** list, select the **vsphere_client** policy, click **Add**, then click **OK** to associate this policy with this resource. For more information, see "Resource Policies Page" in the *Cloud Manager Administrator Reference*.

   If you want to connect multiple vCenter Servers, ensure that you modify the `vsphere.vcenters fact` of the vsphere_client policy as described in the policy comments.

Creating a Policy to Apply to Each Orchestration Resource in the Respective vSphere Environments

You can use separate resources to connect to and manage the different vCenter environments that you have configured. To do this, create a custom policy for each vCenter that you want to manage and assign these policies to the resources that you designate to manage the respective vCenters.

The content of the policy should be similar to the following:

```
<policy>
   <resource>
   <fact name="vsphere.vcenters">
      <array>
       <string>VCENTER1_NAME</string>
    </array>
   </fact>
  </resource>
</policy>
```

In this case, applying this policy along with the vsphere_client.policy to a resource would enable that resource to connect to and manage the vCenter with the name `VCENTER1_NAME`. This name must match the **Account Name** you configured in "Creating Accounts for Each vCenter Environment" on page 90.

11.4.3 Optionally Specifying an Authentication Certificate for Each vCenter Server

The vsphere provisioning adapter job automatically enables a secure SSL connection between your Orchestration Agent and the vCenter Server. This involves some security risk if a malicious user is impersonating your vCenter Server. To avoid this risk, you can explicitly configure the SSL certificate that the Orchestration Agent accepts from the vCenter Server.

We recommend that you review VMware documentation regarding gathering the certificate (http://pubs.vmware.com/vsphere-50/topic/com.vmware.wssdk.dsg.doc_50/sdk_sg_server_certificate_Appendix.6.4.html#991190) used by your vCenter Server's web interface before you proceed further.

When you have gathered the certificate, use the following steps to explicitly configure the certificate:

1 Ensure that the Orchestration Agent is installed and started on a computer in each vCenter environment.

   For more information, see Chapter 5, "Installing Cloud Manager Orchestration Components," on page 39.

   If you are installing the agent to a vSphere environment, you can install the agent either locally on the vCenter Server (the vCenter appliance is not supported), or on a dedicated system (virtual or physical) as long as the OS in that system is supported for the Orchestration Agent.

2 In the Cloud Manager Orchestration Console, log in to the Orchestration Server that you want to use to manage vSphere VMs.

3 In the Explorer tree of the Orchestration Console, select the Orchestration Server or “grid” object, then select the **Authentication** tab in the admin view to open the Authentication page.

4 Create a credential to authenticate to a unique vCenter Server in your enterprise. In most cases, the credential is for “administrator” account of the Windows machine where the vCenter Server is running.

4a On the Authentication page, scroll to the Credential Manager (consisting of the **Stored Credentials** panel and the **Stored Certificates** panel), then click **Add Certificate** to display the **Add Certificate** dialog.

4b Fill in the fields of the dialog:

- **Identifier:** Specify a value in that uniquely identifies the certificate associated with this unique vCenter Server. the identifier should be of the form `vsphere_<YOUR_VCENTER_NAME>`, where `<YOUR_VCENTER_NAME>` is the account name that you configured earlier for the vCenter Server.
- **Location:** Specify the file location of the certificate you gathered previously.
- **Group:** Enter `vsphere` as the group name.

For more information, see “Authentication Page” in the *Cloud Manager Administrator Reference*.

4c Click **Add**.

5 Repeat Step 4 for all of the vCenter Servers you want to connect to.

After you have completed the authentication configuration, continue with “Running Discovery for Multiple Resource Environments” on page 94.

11.5 Discovering vSphere Hosts and VMs

After you configure the vSphere Provisioning Adapter and it is running on the Orchestration Server, you can discover the vSphere hosts and the VMs on them.

- Section 11.5.1, “Running Discovery for a vSphere Environment,” on page 93
- Section 11.5.2, “Running Discovery for Multiple Resource Environments,” on page 94

11.5.1 Running Discovery for a vSphere Environment

Discover the VM images on the vCenter Server and populate the Orchestration Console Explorer tree.

1 From the main menu, select **Provision** > **Select VM Hosts and Repositories** to display the Discover VM Hosts and Repositories dialog.

2 In the Discover VM Hosts and Repositories dialog, select the **vsphere** job, then click **OK**.

TIP: Ensure that this job completes before proceeding to Step 3: Repositories where VMs might reside must be discovered prior to any attempt to discover VM images residing there.

---

3 From the main menu, select **Provision** > **Discover VM Images** to open the Discover VM Images dialog.

4 In the Source Repositories table of the Discover VM Images dialog, select the repositories where vSphere images are stored, click **Add** to move the repositories to the Target Repositories table, then click **OK** to run the image discovery.

11.5.2 Running Discovery for Multiple Resource Environments

Discover the vSphere hosts and repositories in each vCenter Server environment, and then discover the VMs on them to populate the Orchestration Console Explorer tree.

1 From the main menu, select **Provision** > **Select VM Hosts and Repositories** to display the Discover VM Hosts and Repositories dialog.

2 In the Discover VM Hosts and Repositories dialog, select the **vsphere** job, then click **OK**.

When you perform this discovery action, the Orchestration Server runs jobs that discover the VM hosts, repositories, and networks in each of the vSphere environments. On each discovered object, the server also generates a `*.vsphere.vcenter` fact that contains a vCenter ID from the hosting vSphere environment.

After the objects are discovered in the vSphere environments, you can use the Orchestration Server to discover existing VMs in those environments.

TIP: Ensure that this job completes before proceeding. Repositories where VMs might reside must be discovered prior to any attempt to discover VM images residing there.

---

3 From the main menu, select **Provision** > **Discover VM Images** to open the Discover VM Images dialog.

The Orchestration Agent discovers all of the VMs managed in the vSphere environments and places them in the Orchestration model for you to manage.

4 In the Source Repositories table of the Discover VM Images dialog, select the repositories where vSphere images are stored, click **Add** to move the repositories to the Target Repositories table, then click **OK** to run the image discovery.

When a VM with a given name is discovered in two different vSphere environments, the second VM discovered is named in the form of `VMNAME_VCENTERID`, rather than named by appending an incremental number, as explained above. As with other such object names that are automatically generated, these VM names can be changed. See “Resource Object Naming and Renaming” in the *Cloud Manager Administrator Reference*.

12 Configuring vSphereUpdate Client and Monitoring Jobs

Typically, all VM management occurs from within Cloud Manager. The Cloud Manager Orchestration update infrastructure uses the Orchestration vSphereUpdater tool to synchronize events and changes made outside of Cloud Manager in your vSphere vCenter system to the Orchestration Server, which in turn, notifies the Cloud Manager Application Server.

The Cloud Manager Orchestration update infrastructure consists of two main components:

- The vSphereUpdate monitor job, which starts the vSphereUpdate Client component and ensures that it runs when necessary.
- A vSphereUpdate Client component, which is executed by the Orchestration Agent.

For a list of the changes that the update infrastructure currently detects and synchronizes, see "vSphere Server Changes Synchronized to Cloud Manager by vSphereUpdater" in the *Cloud Manager Administrator Reference*.

- Section 12.1, "Understanding vSphereUpdate Monitor Jobs," on page 95
- Section 12.2, "Configuring the vSphereUpdate Client," on page 97
- Section 12.3, "Troubleshooting vSphereUpdater Connections to vCenter," on page 99

12.1 Understanding vSphereUpdate Monitor Jobs

In the Orchestration Server Job Scheduler, the vSphereUpdate monitor job is located in the All jobs group. It is associated with both the `vsphere` provisioning adapter policy (for your vSphere VCenter system configuration information) and the vSphereUpdate policy.

The vSphereUpdate policy specifies the facts defined in Table 12-1. You can modify these facts to accommodate your environment. If you modify the default values, vSphereUpdater uses the preferred values when the job is invoked. Lock the settings to override the default value for all future vSphereUpdate jobs.

For more information about scheduling jobs for the Orchestration Server, see "The Orchestration Server Job Scheduler" in the *Cloud Manager Administrator Reference*.

Table 12-1 *Cluster-Related Facts in the vSphereUpdate Policy*

Fact Name	Type	Description
`jobargs.readTimeout`	Integer	Specifies amount of time (measured in seconds) that calls to the vSphere vCenter system) wait to read data before timing out. The default value is 0 (zero), which will wait infinitely for data. The readTimeout option is mainly used for the WaitForUpdates call, which opens a TCP pipe that waits infinitely until an event is available from the vSphere vCenter system. If the vCenter system goes down or reboots, this call blocks infinitely unless a `readTimeout` fact is set.
`jobargs.max.retry.amounts`	Integer	Specifies the number of consecutive times to try to reconnect from the vSphereUpdater to the vSphere vCenter system if the connection goes down. It attempts to connect every 5 minutes. The default maximum is 3 retry attempts, for a total wait time of 15 minutes. If the connection goes down, vSphereUpdater tries to reconnect to the vSphere vCenter system every five minutes up to the maximum number of attempts. After the reconnection attempts have been exhausted, vSphereUpdater exits its process and logs an error for the connection failure.
`jobargs.zos.proxy.user`	String	Specifies an administrative user used by the Orchestration Console to log in to the Orchestration Server in order to perform update operations there. You must create an administrative user for this purpose, if you have not already done so. The name of this user must contain the word “proxy,” for example, *my_proxy*, or *proxy1*. When you change the value of this fact, you must restart the Orchestration Server. For information about configuring the vSphereUpdate Client, see “Configuring the vSphereUpdate Client” on page 97.
`jobargs.zos.proxy.passwd_validity`	Integer	Specifies the amount of time (measured in seconds) that the `zos.proxy.user` password is valid. **Example:** 86400 (1 day). Although the default value (-1) implies that the password is valid forever, the actual validity time is limited to the uptime of the Orchestration Server. When the password expires, the Orchestration Console is automatically restarted with a new password the next time that the monitor job runs.

Fact Name	Type	Description
`jobargs.debug`	Boolean	Specifies whether you want extra verbose debug logging sent to a job log. **NOTE:** The client logs its output to the `log.txt` and `err.txt` files located in `<agent_install_dir>/node.default/.vSphereUpdate/<host name>/<vcenterId>`.
`jobargs.verbose`	Boolean	Specifies whether you want verbose logging sent to a job log. This fact is implicitly set when `jobargs.debug` is set.
`jobargs.mode`	String	Specifies the type of operation to be executed. An empty string (the default) ensures that a vSphereUpdate monitoring process is running on the configured vSphere clients. There are four other valid values for this string: ◆ **`resetSSL`:** Clears the SSL certificate stored on the client for the Orchestration server. This is a convenience operation that has security implications ◆ **`stop`:** Stops any running vSphereUpdate monitoring processes on the configured vSphere clients. ◆ **`kill`:** Executes an OS level hard stop of the vSphereUpdate monitoring processes. ◆ **`clear`:** Clears the cached Orchestration user connection information, forcing a reconnection to the Orchestration server.
`jobargs.pollInterval`	Integer	Specifies the amount of time (measured in seconds) to wait between each check for repository and network changes in vSphere.

12.2 Configuring the vSphereUpdate Client

1 Create a proxy user:

1a In the Orchestration Console, click **Actions** > **Create User** to open the Create a New User dialog.

1b In the **Source Groups** list, select administrators, then click **Add** to move the administrators user group to the **Target Group** list.

1c In the **New User Name** field, specify a user name, click **Create**, then click **Close**.

This is the proxy user. The user name must contain the word “proxy,” for example, *my_proxy*, or *proxy1*.

2 Modify the vSphereUpdate.policy (or modify the job arguments in the Orchestration Server Job Scheduler) so that `zos.proxy.user` contains the name of the user created in Step 1c:

2a In the Explorer Tree, select the **Policies** group to expand the list of policies included on this grid.

2b Select the **vSphereUpdate** policy to open the Policy Editor view.

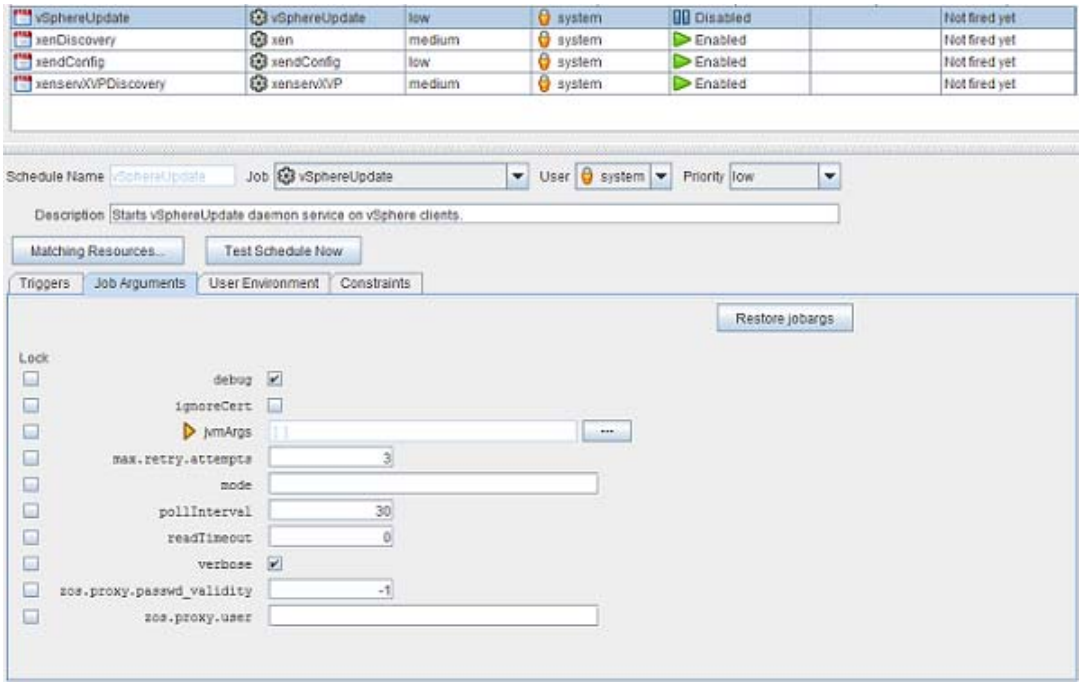

2c Find the `zos.proxy.user` fact in the policy, then specify the name of the proxy user you created in Step 1c as the value for this fact.

2d (Optional) configure other policy settings for the vSphereUpdate Scheduled Job.

See Table 12-1, “Cluster-Related Facts in the vSphereUpdate Policy,” on page 96.

3 Run the vSphereUpdate schedule and job:

3a In the toolbar of the Orchestration Console, select **Scheduler** to open the Orchestration Server Job Scheduler.

3b Select the **vSphereUpdate** schedule, click **Enable**, then click **Run Now**.

4 (Optional) Verify that the update job has run.

4a In the Orchestration Console main menu, select **Jobs** to open the Jobs admin view.

4b In the admin view, locate the VsphereUpdate job that ran last, then select its **Job Log** tab.

You should see something similar to the following in the log:

```
[vrack-vc] checking pid: 5276
[vrack-vc] pid '5276' is still alive
```

The “pid” reference in the log refers to the `javaw.exe` process running on the resource that accesses the vCenter software. You can verify that this process is running in the Windows Task Manager on the VCenter host machine.

12.3 Troubleshooting vSphereUpdater Connections to vCenter

You can improve reliability for the vSphereUpdater connection with VMware vCenter by providing custom values for the read timeout (`readTimeout`) and maximum retry attempts (`max.retry.attempts`) job arguments for vSphereUpdate Scheduled Jobs. These options enable vSphereUpdater to recover after a VMware vCenter crash or reboot.

To set the readTimeout and max.retry.attempts values for vSphereUpdate scheduled jobs:

1. Log in to the Orchestration Server Console.
2. In the Job Scheduler, select the **vSphereUpdate** schedule.
3. Select the **Job Arguments** tab.
4. Specify the **readTimeout** value (in seconds).
5. Specify the **max.retry.attempts** value.
6. (Optional) Lock the options by selecting their Lock check boxes to perpetuate your preferred settings.

13 Configuring Sysprep or Autoprep

When a Cloud Manager provisioning adapter discovers the VMs in your hypervisor, you need to configure those VMs for future provisioning. This section explains sysprep concepts and configuring sysprep for Windows VMs in the Orchestration Console. It also explains autoprep concepts and configuring autoprep for Linux VMs in the Orchestration Console.

- Section 13.1, “Understanding and Configuring Sysprep,” on page 101
- Section 13.2, “Understanding and Configuring Autoprep,” on page 114

13.1 Understanding and Configuring Sysprep

In the Orchestration Console, sysprep refers to the function of preparing unique settings for Windows VMs on VM hosts so that those VMs can be provisioned by the provisioning adapter without creating conflicts and personalizing other information.

As the administrator, you can set facts in the Orchestration Console that can later be automatically applied to a VM clone (by selecting the **Use Autoprep** check box) during a **Provision** or a **Clone** action from a VM template. The sysprep facts can also be manually applied to an existing VM by using the **Personalize** action.

Windows VMs managed by Cloud Manager provisioning adapters can be “sysprepped” in the Orchestration Console. To do so, ensure that you have performed the following prerequisite tasks:

- Create a Windows VM by using the hypervisor.
- Ensure that the VM is discovered by the Orchestration Server.
- Configure the VM (or a template of the VM) by using the information in Section 13.1.2, “Setting Sysprep Facts in the Orchestration Console,” on page 102 and in Section 13.1.3, “Using the Sysprep deploy.cab Files,” on page 110.

When the prerequisites are met, you can proceed to sysprep the VM by using the information in Section 13.1.4, “Applying Sysprep Facts,” on page 112.

This section includes the following information:

- Section 13.1.1, “How Sysprep Works,” on page 101
- Section 13.1.2, “Setting Sysprep Facts in the Orchestration Console,” on page 102
- Section 13.1.3, “Using the Sysprep deploy.cab Files,” on page 110
- Section 13.1.4, “Applying Sysprep Facts,” on page 112
- Section 13.1.5, “Example Sysprep Scenarios,” on page 113
- Section 13.1.6, “Known Sysprep Limitations,” on page 113

13.1.1 How Sysprep Works

Cloud Manager provisioning adapters use the settings specified in the sysprep facts to perform an “unattended mini-setup” to reconfigure the VMs’ Windows guest operating system. The provisioning adapter directly modifies the VM image without need for any agent, so you can configure sysprep on a “cold” VM image in the same way as you can run the **Install Agent** action on that image.

NOTE: You can perform the **Install Agent** and **Personalize** actions at the same time on a Windows VM. The two operations cooperate and complete without conflict.

13.1.2 Setting Sysprep Facts in the Orchestration Console

You can use the Orchestration Console to configure the facts for sysprep of a VM. This section includes information about the Orchestration Console interface where those facts are set.

When you select a Windows VM object in the Explorer tree of the Orchestration Console, click the **Info/Groups** tab to open the Info Groups page, then scroll down to the **Provisioning Information** panel of this page. Open the **Windows Sysprep Config** subpanel and the **Network Sysprep Config** subpanel.

Figure 13-1 *The Sysprep Sections of the Info/Groups Page of a VM Template Object*

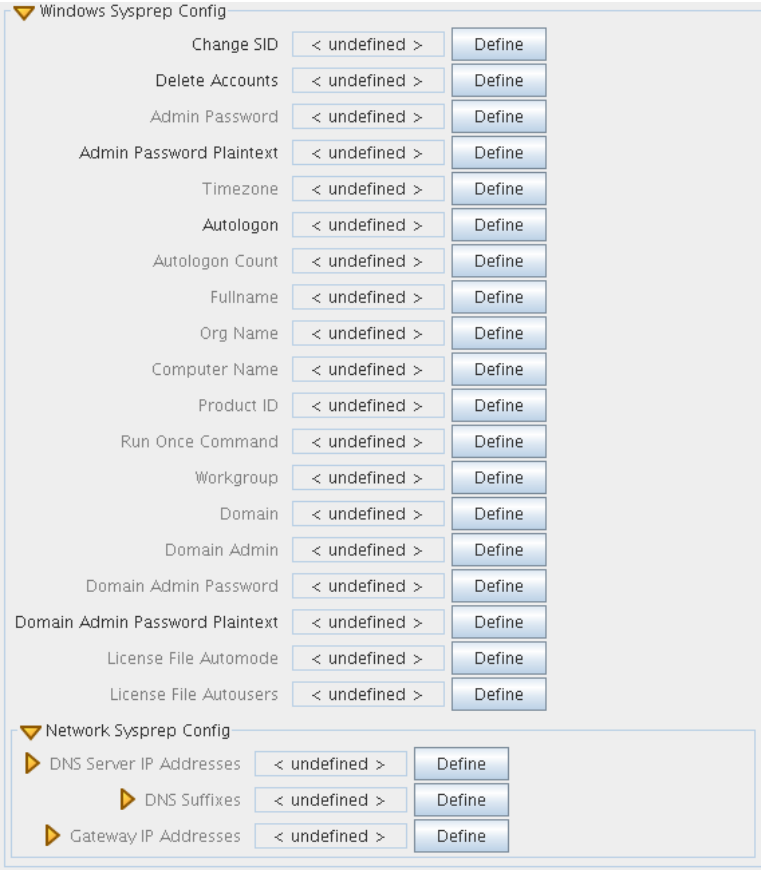

Windows VMs that you clone can be personalized and prepared for provisioning by configuring the facts in this panel. Click **Define** on each field if the value has not been previously configured.

NOTE: You can trigger a sysprep for a Windows VM just by configuring the **Admin Password** field on this page. The provisioning adapter fills in the other required values with reasonable defaults if you don't specify them. For example, the value for the **Computer Name** field uses the VM name in the Orchestration Server by default.

When you finish changing the settings in this panel, right-click the VM and select **Personalize** for the changes to take effect. This sets up a pending customization that does not take place until the VM is powered on (provisioned) again.

IMPORTANT: On VMs managed by any hypervisor, running the **Personalize** action on a templated VM is not supported. Running this action results in failure because it is not supported in the underlying system. When you clone or provision from a templated VM, select the **Use Autoprep** check box.

This section also includes the following information:

- "Sysprep Facts" on page 103
- "Configuring vNIC Sysprep Facts" on page 110

Sysprep Facts

The following table lists the facts that are either required or optional for configuring sysprep.

Table 13-1 *Required or Optional Facts for Sysprep Configuration*

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
Change SID	`<fact name="resource.provisioner.autoprep.options.changeSID" value="false" type="Boolean" />`	Automatic default value provided by the Orchestration Server[1]	The Windows Security ID. If this is marked as true, sysprep generates a new Security ID. If this is not selected, the Orchestration Server defaults the value to true, meaning a new SID is to be generated during sysprep. For newer (that is, unattended `.xml`-based) sysprep, unless this fact is defined and explicitly set to false, the SID is always changed. This is the desired behavior for cloning a Windows machine.

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
Delete Accounts	`<fact name="resource.provisioner.autoprep.options.deleteAccounts" value="false" type="Boolean" />`	Optional[2]	(Windows with vSphere VMs) When this check box is selected (it has a value of true), the Orchestration Server removes all user accounts from the destination VM; the Administrator account and other Windows "standard" accounts are left in place. If it is false, existing accounts from the source VM are retained. No account deletion is done unless this fact is defined and set to true. Also, some versions of Windows sysprep do not support account deletion during sysprep, in which case this flag is ignored.
Admin Password	`<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.value" value="" type="String" />`	Required[3]	The Admin password for this VM. **NOTE:** Only a plaintext admin password is currently supported. You should leave this field blank if **Admin Password Plaintext** is not selected. This fact must be specified or the personalization fails. Windows sysprep requires that this be specified.
Admin Password Plaintext	`<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.plainText" value="false" type="Boolean" />`	Automatic default value provided by the Orchestration Server[1]	When this check box is selected (it has a value of true) the **Admin Password** value is entered in plain text. If not set, this fact defaults to true, indicating that the AdminPassword fact is a plain text password. The Orchestration Server does not support automatic encryption of the password, so if you want to use an encrypted password, you need to know how to encrypt the password correctly for `sysprep.inf`, then enter it as `AdminPassword.value` with this flag set to false.

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
Timezone	`<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.TimeZone" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The time zone of the new VM. For sysprep on Windows versions prior to Vista, (that is, versions using `sysprep.inf`), numeric time zone codes are used. See the Microsoft sysprep documentation for values (for example, 04 indicates PST, 10 indicates MST, 20 indicates Central, 35 indicates EST as defined in the Windows sysprep documentation (http://technet.microsoft.com/en-us/library/cc749073.aspx)). **NOTE:** Ensure that you use the exact text string listed under the Time Zone column heading in the table included in this Microsoft article. For sysprep on Windows Vista and later (that is, versions using `unattend.xml`), full string time zone names are used. Refer to the Microsoft sysprep documentation for the relevant Windows version. If you do not set this fact, the default time zone for `sysprep.inf` (UTC or 85) is used.
Autologon	`<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogon" value="false" type="Boolean" />`	Optional[2]	When this check box is selected (it has a value of true) the VM automatically logs on to the Administrator account by using **AdminPassword**. If the value is false, logon is prompted. If no value is provided, the fact is set to false.
Autologon Count	`<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogon" value="" type="Boolean" />>`	Optional[2]	The limit count for the VM to automatically log on with the Administrator account. **AutoLogon** must be true for this value to be accepted. If a value is not specified for this fact, but **Autologon** is set to true, the value defaults to 1.

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
Fullname	`<fact name="resource.provisioner.autoprep.sysprep.UserData.FullName" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The user's full name required by the Windows OS installer during installation.
Org Name	`<fact name="resource.provisioner.autoprep.sysprep.UserData.OrgName" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The organization name required by the Windows OS installer during installation. This fact is required by sysprep. If the value is not specified, the provided default is Organization Name. This fact is nonessential for Windows functionality.
Computer Name	`<fact name="resource.provisioner.autoprep.sysprep.UserData.ComputerName" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The VM's new host name. If you specify an asterisk (*) in this field, the Orchestration Server generates the name based on the source VM name. This fact is required by sysprep. If the value is not set, the default value is the name of the VM in the Orchestration Server. The name cannot be longer than 15 characters because of a Windows limitation on the length of the computer name. Values longer than 15 characters are not accepted. Because facts can be edited using methods other than the Admin view of the Orchestration Console, be aware that there are other restrictions regarding the characters you can use for the Computer Name. The following characters are not allowed: `whitespace ` ! @ # $ ^ & * () + = [] {} \ | ; : ' " , <> / ?` Other methods you could use to edit the computer name fact might not enforce any restrictions or constraints on the naming. If the naming is invalid, the VM might hang during sysprep.

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
Product ID	`<fact name="resource.provisioner.autoprep.sysprep.UserData.ProductID" value="" type="String" />`	Effectively required[4]	The Windows product key. The ID is obtained from the Windows MSDN CD or from Microsoft. The value is used when building a new VM. This fact is optional for the Orchestration Server, but if the value is not specified, the Windows VM might stop at a user prompt on its console waiting for the entry of the Product ID. Certain versions of Windows, such Windows Server 2008, might not require a product key at installation, but will eventually require it (or a valid license server setup for product activation).
Run Once Command	`<fact name="resource.provisioner.autoprep.sysprep.GuiRunOnce.Command" value="" type="String" />`	Optional[2]	A list of commands separated by new line characters that run the first time a user logs on after the new VM is created. Commands are scheduled by using the `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\RunOnce` registry key. The value does not need to be specified. It is passed to the sysprep answer file only if one or more commands are specified in the list.
Workgroup	`<fact name="resource.provisioner.autoprep.sysprep.Identification.JoinWorkgroup" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The Windows workgroup name. If the VM is joining a domain, use `JoinDomain`. Sysprep requires either a domain or a workgroup to be joined. This fact is ignored if the `Domain` fact and related facts are set, because domain joining takes priority over Workgroup joining. If no domain is being joined, and this fact is not specified, the default value becomes WORKGROUP (default with Windows).

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
Domain	`<fact name="resource.provisioner.autoprep.sysprep.Identification.JoinDomain" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The Windows domain name. If the VM is joining a workgroup, use `JoinWorkgroup`. For joining a domain, `DomainAdmin` and `DomainAdminPassword` must be defined. No default value is provided if this value is not set. Instead, a workgroup is joined with the default name WORKGROUP if no `Workgroup` fact was set.See also: Domain Admin Password (required if a value is set for this fact).
Domain Admin	`<fact name="resource.provisioner.autoprep.sysprep.Identification.DomainAdmin" value="" type="String" />`	Required[3]	Provide a value for this fact when the `Domain` fact has a value. Configuring this fact allows sufficient privileges to the Windows sysprep program to add the new server or workstation to the domain. Normally, this is the Administrator account for the domain. If the `Domain` fact does not have a value, this fact is ignored.
Domain Admin Password	`<fact name="resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.value" value="" type="String" />`	Required[3]	Provide a value for this fact when the `Domain` fact has a value. Configuring this fact allows sufficient privileges to the Windows sysprep program to add the new server or workstation to the domain. Normally, this is the Administrator password for the domain. If the `Domain` fact does not have a value, this fact is ignored.
Domain Admin Password Plaintext	`<fact name="resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.plainText" value="false" type="Boolean" />`	Automatic default value provided by the Orchestration Server[1]	When this check box is selected (it has a value of true), `DomainAdminPassword` is in plaintext. The value defaults to true if it is not set. The value is currently ignored for sysprep, because there are no fields to accept this flag in `sysprep.inf` or `unattend.xml`.

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
License File Automode	`<fact name="resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoMode" value="" type="String" />`	Automatic default value provided by the Orchestration Server[1]	The value in this field is either PerServer or PerSeat. If it is set to PerServer, the `AutoUsers` fact must be set. A value for this fact is required for sysprep on Windows Server products. The provided default is PerServer.
License File Autousers	`<fact name="resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoUsers" value="200" type="Integer" />`	Automatic default value provided by the Orchestration Server[1]	Specify value between 1 and 9999, representing the number of client licenses per seat. A value for this fact is required for sysprep on Windows Server products. The provided default is 5.
Available in Facts tab only (not a string)	`<fact name="resource.provisioner.autoprep.autoprep_complete" value="false" type="Boolean" />`	Automatic default value provided by the Orchestration Server[1]	Under normal circumstances, you should not need to change the value of this fact. For Windows VMs, sysprep should run only when the VM has been cloned from within the Orchestration Server, the very first time that the Personalize action is applied on the new VM. By default, the fact is always false. The Orchestration Server sets the fact to true only when a new clone is created from a template. When the fact is set to true, the system recognizes that `sysprep` has already been applied to the VM. When a personalize action is run thereafter (either from the Orchestration Console or from the Application Console), an entire `sysprep` operation is not run, but the network-related settings for that VM can be changed (that is the IP address -- DHCP or static -- DNS, and gateways). You can manually change the value of this fact back to false if you have a reason to run a full `sysprep` on the VM. Remember that this voids all existing VM personalization.

[1] Facts with automatic default values must be set in the `sysprep.inf` or `unattend.xml` answer file, but the vmprep job provides useful “generic” defaults if any of these facts are not specified. In general, a VM can be successfully personalized using only an Admin Password and Product ID. Some versions of Windows do not require the Product ID if their activation timeout for Windows Product Activation has not yet expired.

[2] Optional facts are never required. They are not placed in `sysprep.inf` or `unattend.xml` if a value is not specified. For the purpose of this documentation, “optional” also means that sysprep continues to function if these facts are not specified; however, optional facts might be needed as part of a Domain join or a similar function.

[3] There is no way to provide default values for required facts, and sysprep fails ungracefully if they are not specified. The vmprep job fails a Personalization action if these facts are not set by the user.

[4] For sysprep, this is the only fact that is “effectively” required. Some versions of Windows can be installed or sysprepped without this value, but most versions stop during sysprep and await user interaction at a Product Key prompt if this value is not specified in `sysprep.inf` or `unattend.xml`.

Configuring vNIC Sysprep Facts

VMs can be prepared for provisioning by configuring the facts in either the **Autoprep Network Adapter** subpanel (Windows VMs) of the vNIC **Info/Groups** panel or the **Sysprep Network Adapter** subpanel (Linux VMs). For more information, see “Defining Autoprep/Sysprep Network Adapter Facts on the vNIC Object” on page 117.

Virtual NIC sysprep facts are always optional. If they are not specified, Windows chooses “typical system” values, including setting vNICS to use DHCP.

13.1.3 Using the Sysprep deploy.cab Files

By default, Microsoft does not include the `sysprep.exe` or `setupcl.exe` utilities needed for sysprep on Windows Server 2003. To provide sysprep functionality for VMs running these Windows versions, the Orchestration Server must have access to compatible `deploy.cab` files from Microsoft. These files can usually be copied directly from the Windows installation CD, or they can be downloaded from Microsoft.

For Windows Vista, Windows Server 2008, and later releases, Microsoft includes the needed sysprep tools on the OS installation and uses a different “answer file” format and utility syntax. For these newer versions of Windows, there is no need to download additional files from Microsoft, or to copy them from the installation CD.

The following instructions apply only if you want to perform sysprep-based personalization on Windows Server 2003 VMs:

- “.cab File Installation Locations” on page 111
- “Detailed Instructions for Downloading .cab Files From Microsoft” on page 111
- “Detailed instructions for Copying deploy.cab from the Windows Installation CD or DVD” on page 112

.cab File Installation Locations

Assuming a normal install of the Orchestration Server, the server's datagrid file tree is located in the `/var/opt/novell/zenworks/zos/server/datagrid/` directory. Copy the `.cab` files to one of the following locations (leaving off the fully qualified portion of the path before "`/datagrid`") as appropriate for the Windows server operating system:

```
dataGrid/files/sysprep/winserver2003_sp1/x86/deploy.cab
dataGrid/files/sysprep/winserver2003_r2/x86/deploy.cab
dataGrid/files/sysprep/winserver2003/x86/deploy.cab
dataGrid/files/sysprep/winserver2003x64/x86_64/deploy.cab
dataGrid/files/sysprep/winserver2003x64_r2/x86_64/deploy.cab
```

Notice that the files are named according to the `resource.os.type` fact, the `resource.os.arch` fact, and (optionally) whether the VM's operating system is SP1, R2, or something similar. The file tree in the list above covers all of the common releases of Windows.The sysprep job looks for the datagrid file in the following path:

`grid:///sysprep/<resource.os.type>_<servicepack>/<resource.os.arch>/deploy.cab`

If the Orchestration Server cannot find the `.cab` file in this path, it looks for the datagrid file in the following path:

`grid:///sysprep/<resource.os.type>/<resource.os.arch>/deploy.cab`

If you want to install the precise version of `deploy.cab` from your Windows CD, use the above convention to copy it to the `/datagrid` directory.

Through testing, NetIQ has determined that only two unique `.cab` files are required to support the most common Windows versions. See the Download Instructions exceptions on the Microsoft site for details. This method works because the Orchestration Server uses only the `sysprep.exe` and `setupcl.exe` executables from the `.cab` files. The other utilities are not used because their purpose is to manually build `sysprep.inf`. the Orchestration Server automatically builds its own answer file as part of the VM personalization process.

Detailed Instructions for Downloading .cab Files From Microsoft

Use the following steps to download the `.cab` files from Microsoft that you need for sysprep.

TIP: You can deploy the `.cab` files using any user account that has admin rights.

To download the Windows Server 2003 `.exe` file containing `deploy.cab`:

1. From a web browser, navigate to the Microsoft Download Center page entitled *System Preparation tool for Windows Server 2003 Service Pack 2 Deployment* (http://www.microsoft.com/downloads/en/details.aspx?FamilyID=93f20bb1-97aa-4356-8b43-9584b7e72556&displaylang=en), then download the Windows Server 2003 sysprep tool, `WindowsServer2003-KB926028-v2-x86-ENU.exe`.
2. Copy the `.exe` file to a suitable location on a Windows physical or virtual machine, then run the executable with the `/x` flag (which specifies file extraction only) to extract `deploy.cab` from the executable bundle.
3. Navigate to the extracted directory where `deploy.cab` is located.
4. Copy `deploy.cab` to the appropriate locations on the Orchestration Server:

VM personalization testing using the two versions of `deploy.cab` files listed above has determined that they are suitable for common versions of Windows. When you have placed copies of the `deploy.cab` file in the proper directories of the Orchestration Server machine, you can perform sysprep personalization on pre-Vista versions of Windows.You can download other, potentially newer, `deploy.cab` files from Microsoft, but be sure you are familiar with how to use the Microsoft sysprep tools and that the version you download matches the version of your VMs. Ensure that you use the file and directory naming conventions explained in this section, so that the personalization system uses the correct `deploy.cab` for the VM being personalized.

Detailed instructions for Copying deploy.cab from the Windows Installation CD or DVD

If the version of `deploy.cab` you download from Microsoft is not suitable for the Windows version on your Windows VMs, you can copy `deploy.cab` for the version of Windows server you need directly from the Microsoft installation CD or DVD. The file is normally located in the following path relative to the CD's root directory:

```
support/tools/deploy.cab
```

Copy `deploy.cab` to the correct location in the datagrid file tree of the Orchestration Server according your Windows version (see ".cab File Installation Locations" on page 111). For example, if your CD is the x86_64 version of Windows Server 2003, copy it to the following location on the Orchestration Server computer:

```
dataGrid/files/sysprep/winserver2003x64_sp2/x86_64/deploy.cab
```

You can also copy `deploy.cab` to the alternate location used for fall back:

```
dataGrid/files/sysprep/winserver2003x64/x86_64/deploy.cab
```

NOTE: If the `.cab` file you download from Microsoft (see "Detailed Instructions for Downloading .cab Files From Microsoft" on page 111) causes problems with sysprep on your VM images, using the method of copying `deploy.cab` described in this section might correct compatibility problems.

13.1.4 Applying Sysprep Facts

The Orchestration Server applies the sysprep facts by launching the vmprep job when the facts are defined. This job runs automatically and applies the appropriate facts to a VM in the following situations:

- When you run a **Personalize** action on any non-templated VM. (See "Right-Click VM Commands" in the *Cloud Manager Procedures Guide*).

  On VMs managed by any hypervisor, running the **Personalize** action on a templated VM is not supported. Running this action results in failure because it is not supported in the underlying system. When you clone or provision from a templated VM, select the **Use Autoprep** check box.

- When you create a VM clone by initiating the **Clone** action on a VM template.

  You must select the **Use Autoprep** check box in the Orchestration Console if autoprep facts are to be used when the **Clone** action is initiated.

You need to ensure that the VM template is set up according to your needs before you clone or provision it, so that the resulting clone meets your needs.

13.1.5 Example Sysprep Scenarios

Scenario 1: You want to create 25 dynamic VM instances to test job provisioning. You will never use these instances again, so you will not personalize them.

You create a VM template by right-clicking a VM, then you select **Create Template**. When the VM Template is created in the Explorer Tree, you define its sysprep facts in the **Info/Groups** page by specifying an asterisk in the **MAC Address** field, then you select the **Use DHCP** check box. This lets the Orchestration Console autogenerate the MAC address and retrieve network data from the DHCP server. For information about setting sysprep facts on each vNIC, see "Virtual NIC Info Panel" in the *Cloud Manager Administrator Reference*.

When the sysprep facts are defined, you provision this template. You right-click the template object and select **Provision**, then in the Provision VM dialog, you specify that you want to provision (create) 25 new VM instances from this template. Provisioning automatically applies the sysprep facts from the template.

Scenario 2: You have created three VM clones in your grid and you want to provision those clones. You want to ensure that the MAC address and other key network information for each clone is unique, even though each clone is a copy of the same OS image. These clones are to be detached later and used for such things as mail servers and web servers. When the clones were first created, sysprep facts were applied, but now you have changed those facts by adding static IP addresses, subnet masks, and gateway addresses for each. Each clone must be "personalized" because of this change to basic network identifiers.

To personalize, you select each Clone object, then define the adapter-specific settings on the **Info/Groups** page by entering IP addresses, subnet masks, and gateway addresses for each adapter. When you have defined the sysprep facts on each VM clone, you right-click each Clone object in turn and select **Personalize** to apply the new network configuration.

For more information, see "Changing a Virtual Machine Template Clone to an Instance" and "Personalize" in the *Cloud Manager Procedures Guide*.

13.1.6 Known Sysprep Limitations

There are some limitations that you need to be aware of when you use sysprep on VMs:

- When you create a template of a VM with the Windows Server 2008 R2 OS, ensure that you configure the sysprep settings for all of its network interfaces using a DHCP connection, not an IP address. This is necessary because of a problem in Windows sysprep that does not remove old IP addresses from the template. The guest OS must not have an IP configured when it is sysprepped.

  When the template has been prepared to use DHCP, subsequent syspreps of the clones of that template can use an IP address.

- Testing has shown that personalizing VMs that have pre-Vista Windows operating systems does not properly configure some network settings. This is a sysprep limitation. The issue is manifest when vNIC-specific settings from sysprep facts on the VM's vNICs are not configured in the VM after personalization.

  To work around this issue, personalize by using DHCP settings. You can do this by leaving the fields blank. DHCP is the default for network settings.

- Sysprep does not work on Windows VMs until you set a value for the `Admin Password` fact: `resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.value`.

  For information about this fact, see **"Admin Password" on page 104**.

- If you clone a Windows Server 2003 VM template originally created in the hypervisor environment, the administrator password for the VM template base image must be blank (no value), or the original VM administrator password is retained and you cannot log in to the cloned VM with the new password.

  Attempting to change an old password value by using the `AdminPassword` entry in `sysprep.inf` does not work, but if the original password value was blank, you can use the `AdminPassword` entry in `sysprep.inf` to provide the password value and log in with that password. The value is applied from the `resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.value` sysprep fact when you select the **Use Autoprep** check box while creating a clone.

- Some sysprep problems have been noted with Windows Product Activation (WPA) functionality, particularly with versions of Windows that require product activation by the end user.

  On Windows Server 2003 SP1, if you have a VM that has passed the initial activation deadline, and you sysprep it, sysprep is applied correctly, but that VM immediately changes to “limited functionality” mode and requires user intervention to reactivate it. Sysprep seems to remove activation and require that the VM be reactivated.

  Further, although there is a Boolean value (AutoActivate) that you can set in the “unattended” section of `sysprep.inf`, setting the value does not always result in auto activation of a VM.

  To avoid this situation, we recommend that you consider volume licensing or similar licensing solutions available from Microsoft that don’t require manual activation by an actual user. For versions of Windows prior to Windows Vista, this would be a VLK (Volume License Key). For Windows Vista or later, Microsoft has license server-based solutions available to handle volume licensing.

In using sysprep on Windows VMs managed by any hypervisor, no special agent or tools are needed for sysprep because of the method used by the Cloud Manager provisioning adapters.

13.2 Understanding and Configuring Autoprep

In the Cloud Manager Orchestration Console, “autoprep” refers to the function of preparing unique network settings for a Linux VM so that VM can be provisioned by its provisioning adapter without creating network conflicts and without customizing other network-related settings.

As the administrator, you can set facts in the Orchestration Console that can later be applied to a VM clone during a **Provision** or a **Clone** action from a VM template. You can also use the **Personalize** action to manually apply autoprep facts to an existing VM.

This section includes the following information:

- Section 13.2.1, “How Autoprep Works,” on page 114
- Section 13.2.2, “Setting Autoprep Facts in the Orchestration Console,” on page 115
- Section 13.2.3, “Applying Autoprep Facts,” on page 118
- Section 13.2.4, “Example Autoprep Scenarios,” on page 119
- Section 13.2.5, “Known Autoprep Limitations,” on page 119

13.2.1 How Autoprep Works

The vmprep job always runs when you clone or provision from a VM template. The job prepares the root disk image of the VM with the defined autoprep settings. On a Linux system, the global autoprep settings for a VM are stored in various configuration files in the `/etc` directory. For example, the host

name is stored in `/etc/HOSTNAME`. Global network properties are stored in `/etc/sysconfig/network/config` and in `/etc/sysconfig/network/dhcp`. Per-NIC properties are written to the various `/etc/sysconfig/network/ifcfg.*` scripts, with one for each virtual NIC.

The vmprep job attempts to identify the disk image with the root partition, then mounts that partition and starts scanning the configuration files to make the necessary changes to the VM configuration file settings.

If the **Use Autoprep** check box in the Orchestration Console is not selected, the vmprep job still runs, but only to change the name of the Orchestration Agent (if installed) on the VM.

If you want a full autoprep with system config changes when cloning or provisioning from template, you need to select **Use Autoprep** check box in the Orchestration Console.

For Linux VMs, autoprep mounts the VM image's root disk image and edits the appropriate files in the `/etc/` directory to make the desired configuration changes. This might include adding network interface configurations to the network configuration scripts. The changes take effect when the VM starts again.

13.2.2 Setting Autoprep Facts in the Orchestration Console

You can use the Orchestration Console to configure the facts for autoprep of a VM. This section includes information about the Orchestration Console interface where those facts are set.

When you select a Linux VM object in the Explorer tree of the Orchestration Console, click the **Info/Groups** tab to open the Info Groups page, then scroll down to the **Provisioning Information** panel of this page. Open the **Linux Autoprep Config** panel and the **Network Autoprep Config** panels.

Figure 13-2 *The Autoprep Sections of the Info/Groups Page of a VM Template Object*

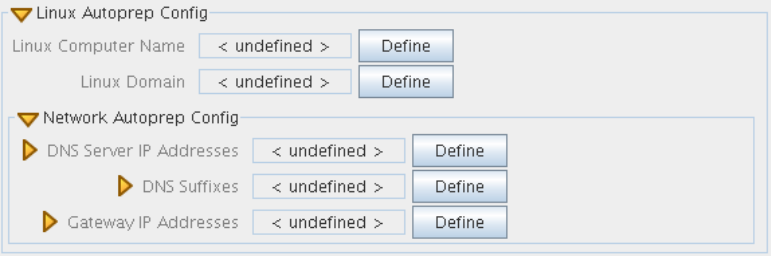

Linux VMs that you clone can be personalized and prepared for provisioning by configuring the facts in this panel. Click **Define** on each field if the value has not been previously configured.

NOTE: When you change any of the settings in this panel, you need to right-click the VM and select **Personalize** for the changes to take effect. This action is in contrast to right-clicking a template, which can apply these settings during a provision or clone operation.

This section also contains this information:

- "Linux Autoprep Config" on page 116
- "Network Autoprep Config" on page 116
- "Defining Autoprep/Sysprep Network Adapter Facts on the vNIC Object" on page 117

Linux Autoprep Config

The settings located in the **Linux Autoprep Config** panel are global to a configuration of a Linux VM and are not specific to a particular network adapter.

NOTE: It is not mandatory to define these facts. If they are left undefined, they are not applied to the "autoprepped" VM.

- **Linux Computer Name:** The network host name of the new VM. If you specify an asterisk (*), the current Grid object ID (`resource.id`) of the new VM is used.

  The Linux Computer Name should be the unqualified computer name without the DNS domain suffix, such as webserver instead of webserver.acme.com.

  In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.linuxglobal.ComputerName`:

  ```
  <fact name="resource.provisioner.autoprep.linuxglobal.ComputerName" value=""
  type="String" />
  ```

- **Linux Domain:** The network domain name where the new VM is a member.

  This field should contain the default DNS domain for the host, such as acme.com.

  In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.linuxglobal.Domain`:

  ```
  <fact name="resource.provisioner.autoprep.linuxglobal.Domain" value=""
  type="String" />
  ```

Network Autoprep Config

This section includes the following fields:

- **DNS Server IP Addresses:** The list of DNS Servers for name for lookup. This setting is only for cloning/personalize actions. For Linux, it should be set only in the VM facts, not in the vNIC facts.

  In the Fact Editor, this fact is listed as an array:

  ```
  <fact name="resource.provisioner.autoprep.DNSServers">
    <array>
      <string></string>
    </array>
  </fact>
  ```

- **DNS Suffixes:** The list of suffixes to append to a name for lookup. This setting is only for cloning/personalize actions. For Linux, it should be set only in the VM facts, not in the vNIC facts.

  ```
  <fact name="resource.provisioner.autoprep.DNSSuffixes">
    <array type="String">
    </array>
  </fact>
  ```

- **Gateway IP Addresses:** The list of Internet gateways available to this VM. This setting is only for cloning/personalize actions. For Linux, it should be set only in the VM facts, not in the vNIC facts.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.provisioner.autoprep.Gateways">
  <array>
    <string></string>
  </array>
</fact>
```

Defining Autoprep/Sysprep Network Adapter Facts on the vNIC Object

VMs can be prepared for provisioning by configuring the facts in either the **Autoprep Network Adapter** subpanel (Windows VMs) of the vNIC **Info/Groups** panel or the **Sysprep Network Adapter** subpanel (Linux VMs). Click **Define** on each field if the value has not been previously configured.

NOTE: When you change any of the settings in this panel, you need to right-click the VM and select **Personalize** for the changes to take effect.

- **MAC Address:** The MAC address of the interface. Specify an asterisk (*) or specify no setting at all to generate a new MAC address. If the value is not set, the existing `vnic.mac` is used.

  IMPORTANT: An unset **MAC Address** fact generates a new MAC address. This is contrary to the current tool tip text.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.MACAddress`:

  ```
  <fact name="vnic.provisioner.autoprep.MACAddress" value="" type="String" />
  ```

- **Use DHCP:** When this check box is selected (it has a value of true), the VM is configured to retrieve its network settings from a DHCP server. If the check box is not selected (it has value of false), you should ensure that the IP address, subnet mask, and gateway address facts are defined. In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.UseDHCP`:

  ```
  <fact name="vnic.provisioner.autoprep.UseDHCP" value="false" type="Boolean" />
  ```

- **IP Address:** The IP address for the adapter.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.IPAddress`:

  ```
  <fact name="vnic.provisioner.autoprep.IPAddress" value="" type="String" />
  ```

- **Subnet Mask:** The subnet mask for this adapter.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.subnetMask`:

  ```
  <fact name="vnic.provisioner.autoprep.subnetMask" value="" type="String" />
  ```

- **Gateway IP Addresses:** (Windows only) A list of the gateway IP addresses available to the interface.

  In the Fact Editor, this fact is listed as an array:

  ```
  <fact name="vnic.provisioner.autoprep.Gateways">
    <array type="String">
    </array>
  </fact>
  ```

  You can edit this array by clicking the ··· button to open an array editor. In this dialog, you can add or remove the IP address or change its order in the array of element choices.

- **DNS from DHCP:** When this check box is selected (it has a value of true), the SUSE VM is configured to retrieve its DNS server settings from DHCP.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.DNSFromDHCP`:

  ```
  <fact name="vnic.provisioner.autoprep.DNSFromDHCP" value="false"
  type="Boolean" />
  ```

- **DNS Server IP Addresses:** (Windows VM only) The adapter's list of DNS servers used for name lookup.

  In the Fact Editor, this fact is listed as an array:

  ```
  <fact name="vnic.provisioner.autoprep.DNSServers">
    <array type="String">
    </array>
  </fact>
  ```

- **DNS Domain:** (Windows VM only) The adapter's DNS domain name.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.DNSDomain`:

  ```
  <fact name="vnic.provisioner.autoprep.DNSDomain" value="" type="String" />
  ```

- **Primary WINS Server:** (Windows VM only) The name of the adapter's primary WINS server.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.primaryWINS`:

  ```
  <fact name="vnic.provisioner.autoprep.primaryWINS" value="" type="String" />
  ```

- **Secondary WINS Server:** (Windows VM only) The name of the adapter's secondary WINS server.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.secondaryWINS`:

  ```
  <fact name="vnic.provisioner.autoprep.secondaryWINS" value="" type="String" />
  ```

- **NetBIOS:** (Windows VM only) The NetBIOS options for this VM. Options include:
  - **EnableNetBIOSviaDhcp**
  - **EnableNetBIOS**
  - **DisableNetBIOS**

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.netBIOS`:

  ```
  <fact name="vnic.provisioner.autoprep.netBIOS" value="" type="String" />
  ```

NOTE: Although you can define individual static settings to be applied to these adapters, autoprep can be useful for provisioning multiple clones with unique, autogenerated MAC addresses and DHCP-defined IP addresses (even though the VM clones are copies of the same VM template OS image) by coupling the autoprep settings on the VM with the autoprep settings on the vNIC object associated with the VM, thus avoiding network conflicts. For more information about vNIC autoprep settings, see "Defining Autoprep/Sysprep Network Adapter Facts on the vNIC Object" on page 117.

13.2.3 Applying Autoprep Facts

The Orchestration Server applies the autoprep facts by launching the vmprep job when the facts are defined. This job runs automatically and applies the appropriate facts to a VM in the following situations:

- When a **Personalize** action is run on any non-template VM. (See "Right-Click VM Commands" in the *Cloud Manager Procedures Guide*).

On VMs managed by any hypervisor, running the **Personalize** action on a templated VM is not supported. Running this action results in failure because it is not supported in the underlying system. When you clone or provision from a templated VM, select the **Use Autoprep** check box.

- When a VM clone is created by initiating the **Clone** action on a VM template.

  Select the **Use Autoprep** check box in the Orchestration Console if autoprep facts are to be used when the **Clone** action is initiated.

- When a VM clone is created by initiating a **Provision** action on a VM template.

  Select the **Use Autoprep** check box in the Orchestration Console if autoprep facts are to be used when the **Clone** action is initiated.

13.2.4 Example Autoprep Scenarios

Scenario 1: You want to create 25 dynamic VM instances to test job provisioning. You will never use these instances again.

You create a VM template by right-clicking a VM, then you select **Create Template**. When the VM Template is created in the Explorer Tree, you define its autoprep facts in the **Info/Groups** page of the vNIC object by specifying an asterisk in the **MAC Address** field, then you select the **Use DHCP** check box. This lets the Orchestration Console autogenerate the MAC address and retrieve network data from the DHCP server. For information about setting autoprep facts on each vNIC, see "Virtual NIC Info Panel" in the *Cloud Manager Administrator Reference*.

When the autoprep facts are defined, you provision this template. You right-click the template object and select **Provision**, then in the Provision VM dialog, you specify that you want to provision (create) 25 new VM instances from this template. Provisioning automatically applies the autoprep facts from the template if the **Use Autoprep** check box is selected.

Scenario 2: You have created three VM clones in your grid and you want to provision those clones. You want to ensure that the MAC address and other key network information for each clone is unique, even though each clone is a copy of the same OS image. These clones are to be detached later and used for such things as mail servers and web servers. When the clones were first created, autoprep facts were applied, but now you have changed those facts by adding static IP addresses, subnet masks, and gateway addresses for each. Each clone must be "personalized" because of this change to basic network identifiers.

To personalize, you select each Clone object, then define the adapter-specific settings on the **Info/Groups** page of each of the VM's vNICs by entering IP addresses, DNS suffixes, and gateway addresses for each vNIC in the **Network Autoprep Config** subpanel. When you have defined the autoprep facts on each VM clone, you right-click each Clone object in turn and select **Personalize** to apply the new network configuration.

For more information, see "Changing a Virtual Machine Template Clone to an Instance" and "Personalize" in the *Cloud Manager Procedures Guide*.

13.2.5 Known Autoprep Limitations

There are some limitations that you need to be aware of when you use autoprep:

- Currently, the **Gateway IP Addresses** setting in the **Info/Groups tab** for a VM object is available in a list box.

Because the Linux VM OS accepts only one default gateway, it accepts only the first setting in the list as the actual gateway IP address. The other settings are ignored.

- Cloud Manager provisioning adapters check the setting for the host name. If the host name is not set, the setting defaults to the VM name in the Orchestration Server.

14 Using the Cloud Manager Application Server Configuration Tool

After you have installed the Cloud Manager components, you need to configure the system according to your data center environment architecture and your objectives for using the product. Cloud Manager provides a configuration tool to help you.

The Cloud Manager configuration tool is highly interactive, detecting your Cloud Manager server operating system and its present configuration. It also detects the installation of the Cloud Manager Orchestration components on the server and gives you the opportunity to configure them by running a separate but related tool.

TIP: In a production environment, we recommend running the Orchestration configuration tool and the Cloud Manager Application configuration tool on different servers.

For more information about the Orchestration configuration tools, see Chapter 7, “Configuring Cloud Manager Orchestration Components,” on page 59.

The Cloud Manager configuration tool also gives you the option to install the product in a demonstration mode, building all of the components (including an embedded ApacheDS LDAP with default users already set up) that you need if you want to demonstrate or conceptualize product functionality. You should not use this “demo mode” if you have installed Cloud Manager in your production environment and you want to configure it there.

This section of the documentation does not discuss the concepts of the demo mode (how to run it or what to observe in it).

The configuration tool includes a script with several segments. Each segment prompts for information and then executes the configuration with the information you provide. The following sections provide the detail about the segments of the script:

- Section 14.1, “Configuring the PostgreSQL Database Connection and Credentials,” on page 122
- Section 14.2, “Configuring Auto-Vacuum Settings for Tables in the PostgreSQL Database,” on page 125
- Section 14.3, “Configuring Cloud Manager to Use Authentication Sources,” on page 126
- Section 14.4, “Installing and Configuring Other Cloud Manager Feature Settings,” on page 130

14.1 Configuring the PostgreSQL Database Connection and Credentials

The Cloud Manager installation pattern includes a `postgresql-server` package. This package can be installed with Cloud Manager on the local host by default. No matter when it is installed, however, a PostgreSQL ORDBMS is required for Cloud Manager. This product uses a dedicated database in PostgreSQL to store all of its data.

This section helps you to prepare the information you need to configure the PostgreSQL instance you use for Cloud Manager.

1 Ensure that you know the information you are prompted to provide during the PostgreSQL configuration:

Information Needed for Configuration	Description
Database server	You need to know the PostgreSQL database server host name or IP address. Unless you chose not to install the `postgres` package during the install, the Cloud Manager Application Server installs the packages in this pattern on the same server where you installed Cloud Manager. The default is `localhost`.
Autoconfigure an unconfigured PostgreSQL installation?	If you install a PostgreSQL ORDBMS intended for Cloud Manager but you have not yet configured it, Cloud Manager can autoconfigure it for your environment. Autoconfigure sets up the PostgreSQL authentication method, changes the default `postgres` user password, and configures the database for local connections only. If you choose to autoconfigure, the tool sets up the environment and exits without displaying a summary of the settings it used. If you want to troubleshoot database problems, you can view the settings. Autoconfigure does the following: ◆ Generates a `postgres` user password and copies it to `/etc/opt/netiq/cloudmanager/etc/pgusr.in`. ◆ Creates a database and names it `cloudmanager`. ◆ Creates a database user and names it `cmadmin`. ◆ Generates a database password for `cmadmin` and copies it to `/etc/opt/netiq/cloudmanager/etc/com.novell.ncm.backend.connpool.cfg` If you choose not to autoconfigure because your PostgreSQL instance is already configured or because it is remotely located, the tool directs you to supply information for a new database configuration for Cloud Manager. **IMPORTANT:** Running autoconfigure on a PostgreSQL instance that is already configured causes autoconfigure to fail.

Information Needed for Configuration	Description
Database server port	You need to know the port that your PostgreSQL server uses for outside communication. The default is `5432`. **NOTE:** This documentation does not discuss the configuration for the PostgreSQL server.
Create a new PostgreSQL database? (You can choose to use an existing database instead of creating an new one.)	If you want to use Cloud Manager with a fresh database, you can choose to create that database. The configuration tool configures that database with data that you supply when prompted: ◆ **Administrator Name:** An administrative user with permission to create a database. ◆ **Database Administrator Password:** The password that the Administrator designated above uses to log in to the database. ◆ **Database Name:** An arbitrary name you specify to identify the database. The default is `cloudmanager`. ◆ **Database User Name:** A user to be created who must have read/write permissions to the database. The default is `cmadmin`. ◆ **Database User Password:** The password that the database user (designated above) uses to log in to the database.
Use an existing PostgreSQL database	If you want to use an existing PostgreSQL database, that database should not have been used prior for any other purpose. The configuration tool configures that database with authentication data that you supply when prompted: ◆ **Database Name:** An arbitrary name you specify to identify the database. The default is `cloudmanager`. ◆ **Database User Name:** A user to be created who must have read/write permissions to the database. ◆ **Database User Password:** The password that the database user (designated above) uses to log in to the database.

2 At the Cloud Manager Application Server, run the Cloud Manager Application configuration tool:

```
/opt/netiq/cloudmanager/configurator/config
```

The tool displays the first segment of its configuration script:

```
Welcome to the Cloud Manager configuration utility.

INSTALLATION OPTIONS MENU

Select products to configure

 #   selected  Item
 1)     no     NetIQ Cloud Manager - Server
 2)     no     NetIQ Cloud Manager - Manage Authentication
 3)     no     NetIQ Cloud Manager - Manage Certificates

Select from the following:
  1 - 3)  toggle selection status
       a)  all
       n)  none
       f)  finished making selections
       q)  quit -- exit the program
Selection [f]:
```

3 Specify `1` to configure the Cloud Manager Application Server, then enter `f` to finish the selection and move to the **PostgreSQL Database Connection** segment of the script.

```
POSTGRESQL DATABASE CONNECTION

This segment of the configuration utility lets you provide PostgreSQL
authentication information to be used by NetIQ Cloud Manager (NCM).

If you want to install Postgres to a local database and Postgres has not been
configured, you can choose to configure Postgres automatically.

If you choose to install to an existing Postgres server, you need the
following information:

    -  The Postgres server IP Address and the port where the service is running
    -  A username with permission to create the NCM database and user

       or

       The database name you want to populate, along with a username with write
       permission to that database.

Press <RETURN> to continue...
```

4 Follow the prompts to complete the PostgreSQL configuration. Use the information you collected in Step 1 on page 122 as the script prompts you.

The configuration tool checks the database server and the database instance you specify, using the newly-defined credentials to ensure that the database instance and the database user can be created.

Following the PostgreSQL configuration, continue with Chapter 14.3, “Configuring Cloud Manager to Use Authentication Sources,” on page 126 to configure the authentication sources you want to use with Cloud Manager.

14.2 Configuring Auto-Vacuum Settings for Tables in the PostgreSQL Database

In a large-scale datacenter, the tables in the Application Server PostgreSQL database can grow quite large. Performance can degrade significantly if stale and temporary data are not systematically removed. Vacuuming cleans up stale or temporary data in a table, and analyzing refreshes its knowledge of all the tables for the query planner.

PostgreSQL database tables are auto-vacuumed by default when 20% of the rows plus 50 rows are inserted, updated, or deleted. Tables are auto-analyzed when a threshold is met for 10% of the rows plus 50 rows. For example, a table with 10000 rows is not auto-vacuumed until 2050 rows are inserted, updated, or deleted. That same table is auto-analyzed when 1050 rows are inserted, updated, or deleted.

The default auto-vacuum analyze and vacuum settings are sufficient for a small deployment, but the percentage thresholds take longer to trigger as the tables grow larger. Performance degrades significantly before the auto-vacuum vacuuming and analyzing occurs.

NOTE: For larger installations (that is, thousands of workloads), we recommend that you customize the settings as follows on individual tables to ensure optimal performance for data access:

- The scale factor should be set to zero for both vacuum and analyze auto-vacuum settings.
- The threshold should be set to 1000 for both the vacuum and analyze threshold settings.

The following PostgreSQL statements will ensure that auto-vacuum vacuum and analyze are run every 1000 updates (inserts, modifies, or deletes) on the tables that are usually the biggest in the Cloud Manager Application Server installation.

```
ALTER TABLE novell.workload_cost SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.workload_cost SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.workload_cost SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.workload_cost SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE novell.cm_domain_obj SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.cm_domain_obj SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.cm_domain_obj SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.cm_domain_obj SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE novell.pso_object SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.pso_object SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.pso_object SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.pso_object SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE novell.disk_data SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.disk_data SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.disk_data SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.disk_data SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE wl_disk_data SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE wl_disk_data SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE wl_disk_data SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE wl_disk_data SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE novell.nic_data SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.nic_data SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.nic_data SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.nic_data SET (autovacuum_analyze_threshold = 1000);
```

```
ALTER TABLE novell.wl_nic_data SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.wl_nic_data SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.wl_nic_data SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.wl_nic_data SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE novell.workload SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.workload SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.workload SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.workload SET (autovacuum_analyze_threshold = 1000);

ALTER TABLE novell.business_service SET (autovacuum_vacuum_scale_factor = 0.0);
ALTER TABLE novell.business_service SET (autovacuum_vacuum_threshold = 1000);
ALTER TABLE novell.business_service SET (autovacuum_analyze_scale_factor = 0.0);
ALTER TABLE novell.business_service SET (autovacuum_analyze_threshold = 1000);
```

TIP: The following query provides information about the last time the Novell schema tables were vacuumed or analyzed:

```
select relname,last_vacuum, last_autovacuum, last_analyze, vacuum_count,
autovacuum_count,
  last_autoanalyze from pg_stat_user_tables where schemaname = 'novell' order by
relname ASC;
```

14.3 Configuring Cloud Manager to Use Authentication Sources

The instructions in this section assume that you have already used the configuration tool to configure the PostgreSQL database use by the Cloud Manager Application Server, as described in Chapter 14.1, “Configuring the PostgreSQL Database Connection and Credentials,” on page 122.

The Net IQ Cloud Manager Application Server can connect to and search several different kinds of authentication sources to collect information about users in those sources. These are the users that can be authorized, depending on their individual roles, to log into Cloud Manager as Cloud Manager users.

The Cloud Manager Application Server configuration tool includes a segment that displays directly after the PostgreSQL Configuration segment of the script, prompting you to choose an authentication source and asking for specific information that allows Cloud Manager connection to that source.

NOTE: If you configured authentication sources in a previous configuration session, you can manage those configuration settings in a new session. The tool provides a new option (`Cloud Manager - Manage Authentication`) that you can select to make authentication configuration changes subsequent to your initial work.

This section discusses the authentication source options in Cloud Manager and how to obtain the data you provide for the tool. The section also includes an explanation of the setup you need to perform, if any, to prepare each of these authentication sources for connection to Cloud Manager.

- Section 14.3.1, “Configuring Authentication to an LDAP Directory,” on page 127
- Section 14.3.2, “Configuring Authentication to NetIQ Access Manager,” on page 129

14.3.1 Configuring Authentication to an LDAP Directory

The Cloud Manager administrator can choose to authenticate users through a supported Lightweight Directory Access Protocol (LDAP) directory service, either Microsoft Active Directory or NetIQ eDirectory. Cloud Manager users must have an account in the LDAP directory and must be members of the Cloud Manager user group. In addition, the LDAP user you specify as the read-only user must have All Attribute access to the area of the directory to be used by Cloud Manager.

You can also choose to add the Secure Sockets Layer (SSL) protocol to manage the security of authentication data being passed between Cloud Manager and LDAP. Adding SSL to the authentication process adds encryption and verification the process.

This section helps you to prepare the information you need to configure LDAP for Cloud Manager authentication. If you want to use another authentication service, see Section 14.3.2, “Configuring Authentication to NetIQ Access Manager,” on page 129.

1 Ensure that you know the information you are prompted to provide during the LDAP configuration:

Information Needed for LDAP Configuration	Description
Do you want to use SSL with LDAP?	If you respond with “yes” to this question, you are asked for an SSL certificate later in the configuration.
LDAP Source	You need to select the LDAP source for use with Cloud Manager, either NetIQ eDirectory or Microsoft Active Directory.
LDAP host address	This is the address (DNS name or IP address) of the LDAP host that Cloud Manager can connect to for authentication. If you chose to use SSL with LDAP, this address should match the subject of the certificate issued for the LDAP host. The configuration tool immediately validates this address when you specify it.
LDAP port	Designate the port where you want the LDAP server to listen for communication from Cloud Manager. If you are using SSL, the default port is 636. If you chose not to use SSL, the default port is 389.
Path to SSL certificate on LDAP server	This is the file system path to the SSL certificate you previously copied to the LDAP server. The certificate must be in DER format.[1] You need to use this setting only if you want to use SSL with the LDAP authentication.
LDAP read-only user DN	Specify the distinguished name (DN) of an existing LDAP read-only user who has read access to the LDAP directory. This user must have All Attribute read rights to the area of the directory that is to be used for Cloud Manager.
LDAP read-only user’s password	Specify the password for the LDAP read-only user. When you specify the user password, the configuration tool immediately attempts an SSL authentication to validate the existence of this user and password.

Information Needed for LDAP Configuration	Description
Cloud Manager LDAP user DN	Specify the DN of an existing LDAP user whom you want to designate as the Cloud Manager administrator. When you specify this LDAP user, the configuration tool immediately attempts to locate the user in LDAP, then asks you to verify that this is the user you want to designate as the Cloud Manager administrator. Ensure that the `mail` attribute is set for this user in LDAP.
LDAP DN of NCM Users	Specify the DN of the LDAP container where the users whom you want to log in to Cloud Manager already exist. This is the parent context of users that will be allowed to log in to the Cloud Manager Application Console. All subdirectories and users are included by default. Ensure that all users, regardless of their context in this container, have their email domain configured prior to logging into the Application Console. **NOTE:** You can use the Cloud Manager Application Console later to import users who do not currently exist in this DN.

[1] Use the following command on a Linux machine to fetch the certificate and then copy it to another machine if needed.

```
echo 'GET / 1.0' | openssl s_client -connect <server_ip_addr_or_dns>:<port>|
sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' >ldap.pem
```

The following command converts the certificate to DER format (required by Cloud Manager):

```
openssl x509 -in ldap.pem -inform PEM -out ldap.cer -outform DER
```

2. Continue running the configuration tool (`/opt/netiq/cloudmanager/configurator/config`). In the configuration segment following the configuration of the PostgreSQL database, the tool displays the following text:

```
Authentication Type

1) LDAP
2) NAM

Selection:
```

3. Specify `1` (LDAP) as the authentication type you want to configure.
4. Follow the prompts and use the information you gathered in Step 1 to complete this segment of the configuration.

After the LDAP authentication configuration, continue with Chapter 14.4, “Installing and Configuring Other Cloud Manager Feature Settings,” on page 130.

14.3.2 Configuring Authentication to NetIQ Access Manager

The Cloud Manager administrator can choose to authenticate customers through NetIQ Access Manager (NAM).

This section helps you to prepare the information you need to configure Cloud Manager authentication through NetIQ Access Manager. If you want to use some other authentication service, see Section 14.3.1, “Configuring Authentication to an LDAP Directory,” on page 127.

If you want to learn more about NetIQ Access Manager, see the NetIQ Access Manager Documentation website (https://www.netiq.com/documentation/access-manager/).

1 Ensure that you know the information you are prompted to provide during the Access Manager authentication configuration:

Information Needed to Configure Authentication to NAM	Description
Cloud Manager Administrator user name	Specify the initial user name that you want to designate as the Cloud Manager administrator. This should be the new administrator’s login name or Common Name (CN) and must already exist in your LDAP directory.
Cloud Manager Administrator email address	Specify the email address of the user you want to be the Cloud Manager administrator. This email address must already exist as an LDAP attribute of the future administrator. If the user has more than one email address, use the first address in the email attributes list. Cloud Manager uses this email address to determine the administrative permissions to apply to the user.

As you continue running the configuration tool (`/opt/netiq/cloudmanager/configurator/config`) following the configuration of the PostgreSQL database, the tool displays the following text:

```
Authentication Type

1) LDAP
2) NAM

Selection:
```

2 Specify `2` (NAM) as the Authentication Type you want to configure.

3 Follow the prompts and use the information you gathered in Step 1 on page 129 to complete this segment of the configuration.

After the NetIQ Access Manager authentication configuration, continue with Chapter 14.4, “Installing and Configuring Other Cloud Manager Feature Settings,” on page 130.

14.4 Installing and Configuring Other Cloud Manager Feature Settings

After you have completed configuring the Cloud Manager Application Server to use your chosen configuration source, you must use the Cloud Manager configuration tool to install or configure other Cloud Manager features that help you administer Cloud Manager.

- Section 14.4.1, “Installing the Cloud Manager Application Console,” on page 130
- Section 14.4.2, “Configuring the Cloud Manager Web Server (Jetty),” on page 130
- Section 14.4.3, “Configuring the Cloud Manager Web Server to Use SSL,” on page 131
- Section 14.4.4, “Configuring Cloud Manager SMTP Mail Settings,” on page 131
- Section 14.4.5, “Configuring Cloud Manager System Shell Login Information,” on page 132

14.4.1 Installing the Cloud Manager Application Console

The first feature that the configuration tool can install is the Cloud Manager Application Console. The console is a web-based user interface that lets you manage your Cloud Manager system. The console’s display layout varies, depending on the role of the user who logs in. We recommend that you install this UI when prompted, unless you choose to create your own customer Web UI.

For more information, see the *Cloud Manager Procedures Guide*.

14.4.2 Configuring the Cloud Manager Web Server (Jetty)

The Cloud Manager configuration tool lets you decide whether to integrate SSL with the Cloud Manager Web server (Jetty). If you want you configure a secure connection between Cloud Manager Orchestration Server and the Cloud Manager Application Server, you need to answer “yes” to the following question:

```
Choose whether to configure the Cloud Manager web server to use SSL.

Use SSL with Jetty? (yes/no):
```

If you choose to use SSL, ensure that you know the information are prompted to provide during the Jetty SSL configuration:

Information Needed to Configure SSL Use with Jetty	Description
Web Console HTTPS Port	Specify the secure port for the Cloud Manager Application Console. By default, this is port 8183, but you can specify any unused secure port.
Web Console HTTP Port	Specify the HTTP (non-secure) port for the Cloud Manager Application Console. If you chose to enable SSL for Jetty, Cloud Manager disables this port in `jetty.xml` for security purposes. You can re-enable the port by uncommenting the relevant section of the file.

14.4.3 Configuring the Cloud Manager Web Server to Use SSL

If you choose to use SSL with Cloud Manager's Jetty Web server, you need to provide Secure Socket Layer (SSL) information that the Cloud Manager Application Server can use to provide a secure connection.

When the configuration tool displays its SSL configuration segment, it immediately detects the existing DNS name of the server where you are performing the configuration. Because this DNS name must match the subject of the security certificate, you can change the DNS name to match the subject of an existing certificate.

The configuration tool lets you choose to use either a self-signed certificate generated by the server, or an existing certificate that you can import. The configuration is based on the details you provide after that initial determination:

```
Select 'yes' if you want to use an existing certificate for <detected
_dns_hostname>. If you select 'no', Cloud Manager will use a self-signed
certificate.

Use existing certificate? (y/n):
```

Ensure that you are prepared with the following information you are prompted to provide for configuring the Cloud Manager Web Server to use an imported SSL certificate:

Information Needed to Configure an Imported SSL Certificate	Description
Path to the Cloud Manager Server Certificate	Specify the path to an existing public certificate (in PEM format) that you want to import and use on this server. For example: `/home/jdoe/cloudmgr/newcert.pem`
Path to the Cloud Manager Server Private Key	Specify the path to the private key file of this server. This must be the private key file (in PEM format) that is provided by your trusted certificate authority. For example: `/home/jdoe/cloudmgr/newkey.pem`
Private Keystore Password	Specify the password you want to use for decrypting the private key file exclusively for Cloud Manager. If you don't want to use a password, press Enter when the tool prompts you with this question.

14.4.4 Configuring Cloud Manager SMTP Mail Settings

Cloud Manager uses SMTP messaging to send notifications about pending or completed system tasks and Business Service status. These notifications are sent from a system-like user account to a Cloud Manager user who receives a preconfigured message appropriate for his or her role and based on conditions or events occurring in the Cloud Manager system.

The Cloud Manager configuration tool lets you decide whether to configure mail settings for the system.

If you choose to use email in this way, you need to answer "yes" to the following question:

```
Configure the SMTP mail settings at this time? (yes/no):
```

If you choose to use e-mail, ensure that you know the information you are prompted to provide during the email configuration segment of the configuration:

Information Needed to Configure SMTP Mail Settings	Description
Email Address of Message Source	Specify the email address from which all system notifications are to be sent. This should be a “no-reply” address because the message is automatically generated from the Cloud Manager system.
Cloud Manager SMTP Host	Specify the DNS name of the SMTP host you want to use with Cloud Manager, for example: `smtp.example.test`.
SMTP Port	Specify the port that the SMTP server is listening on. The default setting is port 25, but you can specify another port if you want to.

If your SMTP server requires authentication, you can configure SMTP later in the Cloud Manager Application Console.

14.4.5 Configuring Cloud Manager System Shell Login Information

As the system administrator, you have access to the inner workings of Cloud Manager. You can access the system through an Apache Karaf shell or through the Karaf Web Console (`http://<cloud_manager_server_address>:8181/system/console/bundles`). This segment of the configuration tool process lets you establish the login credentials for the Karaf system administrator.

The credentials you are prompted to provide for the system administrator configuration are independent of any other credentials for the Cloud Manager System.

Information Needed to Configure Authentication for the System Shell	Description
System User	Specify the initial user name that you want to designate as the Karaf system user.
System User’s Password	Specify the password of the system user. This does not need to correlate to any directory password. It is stored in the `users.properties` file located in `/etc/opt/netiq/cloudmanager/etc`.

IV Advanced Installation and Integration Topics

This section includes installation information for advanced Cloud Manager administrators.

- Chapter 15, “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support,” on page 135
- Chapter 16, “Configuring the Orchestration Server to Use an Audit Database,” on page 149
- Chapter 17, “Integrating the Orchestration Server with a Sentinel Collector,” on page 161
- Chapter 18, “Configuring Secure Authentication Sources to Communicate with Cloud Manager,” on page 171

15 Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support

Ensuring maximum service-level availability and data protection is paramount to enterprise IT infrastructure. Automated failure detection and recovery prevents downtime, and reduces the financial and operational impact of outages to the business. Highly available infrastructure is a key requirement for IT decision makers.

The Orchestration Server is a critical component of your enterprise infrastructure. It continuously monitors and manages physical servers and virtual machines (VMs), and provides high availability for virtual machines by automatically restarting them on other physical servers if the server they are running on becomes unavailable because of a planned or unplanned outage. Therefore, the Orchestration Server itself must be highly available.

This guide describes how to configure the Cloud Manager Orchestration Server in a high availability SUSE Linux cluster and how to provide both service-level restart for the Orchestration Server and failover among the physical servers of a SUSE Linux cluster to ensure that the server remains available and responsive to the infrastructure that it manages.

- Section 15.1, “Overview,” on page 135
- Section 15.2, “Orchestration Server Failover Behaviors,” on page 136
- Section 15.3, “Installing the Orchestration Server to a SLE HAE Cluster Environment,” on page 137
- Section 15.4, “Configuring the Orchestration Server for High Availability,” on page 143
- Section 15.5, “Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster,” on page 147
- Section 15.6, “Creating the Server Cluster Resource Group,” on page 147
- Section 15.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 148
- Section 15.8, “Installing and Configuring other Orchestration Components to the High Availability Grid,” on page 148

15.1 Overview

The following figure illustrates how the Orchestration Server is configured for use in a high availability environment.

Figure 15-1 *The Orchestration Server in a Clustered, High Availability Environment*

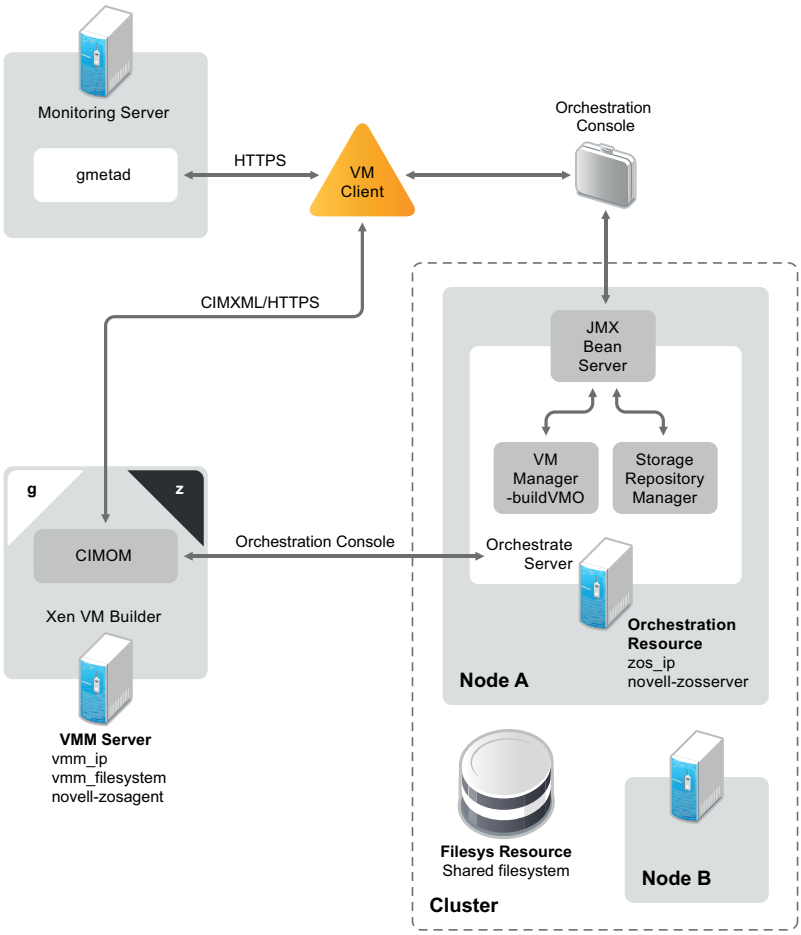

15.2 Orchestration Server Failover Behaviors

This section includes information to help you understand the failover behavior of the Orchestration Server in a high availability environment.

- Section 15.2.1, “Use Case 1: Orchestration Server Failover,” on page 136
- Section 15.2.2, “Use Case 2: VM Builder Behavior at Orchestration Server Failover and Failback,” on page 137
- Section 15.2.3, “Use Case 3: Monitoring Behavior at Orchestration Server Failover and Failback,” on page 137

15.2.1 Use Case 1: Orchestration Server Failover

If the primary node in the Orchestration Server cluster fails, you should see a job restart on another Orchestration Server in the cluster. The job must have been flagged as restartable. For more information, see Section 15.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 148.

When the Orchestration Server fails over to a new node, the Orchestration Agents reauthenticate with the new Orchestration Server instance.

15.2.2 Use Case 2: VM Builder Behavior at Orchestration Server Failover and Failback

If the Orchestration Server fails, any VM builds in progress are canceled and potentially incomplete residual artifacts of the build are cleaned up. When the Orchestration Server restarts or when it fails over in the cluster (the server operates identically in these scenarios), select jobs are run to determine the state of the grid. If the grid was set up with an audit database, the job running at failure time shows as canceled.

15.2.3 Use Case 3: Monitoring Behavior at Orchestration Server Failover and Failback

The Cloud Manager Monitoring Server and the Monitoring Agent are installed on the same server as the Orchestration Server in the high availability Orchestration Server cluster. The Monitoring Agent reports data to the Cloud Manager Monitoring Server.

The Monitoring Server and the Monitoring Agent services are made highly available along with the Orchestration Server and move between clustered machines as the Orchestration Server does. If a monitoring agent is installed on an Orchestration Server and if that server goes down, the server is displayed as “Down” in the Cloud Manager VM Client (Monitoring view).

15.3 Installing the Orchestration Server to a SLE HAE Cluster Environment

This section includes information to help you install Orchestration Server components in a SUSE Linux Enterprise (SLE) High Availability Extension (HAE) environment. The sequence below is the supported method for configuring this environment.

1. Section 15.3.1, “Meeting the Prerequisites,” on page 137
2. Section 15.3.2, “Installing the SLE High Availability Extension,” on page 139
3. Section 15.3.3, “Configuring Nodes with Time Synchronization and Installing Pacemaker to Each Node,” on page 139
4. Section 15.3.4, “Setting Up the OCFS2 File System for High Availability,” on page 141
5. Section 15.3.5, “Installing the Orchestration Server on the First Cluster Node,” on page 142

NOTE: Upgrading from earlier versions of Cloud Manager Orchestration to a high availability environment is supported. For more information, see Chapter 22, “Upgrading Orchestration Server Components,” on page 189.

15.3.1 Meeting the Prerequisites

The environment where the Orchestration Server is installed must meet the operating system, hardware, and software requirements for high availability. This section includes the following information to help you understand those requirements.

- “Hardware Requirements for Creating a High Availability Environment” on page 138
- “Software Requirements for Creating a High Availability Environment” on page 138
- “Other Requirements” on page 139

Hardware Requirements for Creating a High Availability Environment

Ensure that the following hardware requirements are met for creating a high availability environment for the Orchestration Server:

- A minimum of two physical servers configured with the same hardware. NetIQ Corporation recommends a two-node or three-node configuration. These servers are the nodes of the cluster where Orchestration Server is installed and are a key part of the high availability infrastructure.
- Each server must have dual network interface cards (NICs), supporting a minimum data rate of 100 Mbps Ethernet communications.
- A Fibre Channel or iSCSI Storage Area Network (SAN) or network storage, with a shared disk connected to each of the servers that you plan to use in the cluster.
- A STONITH (shoot the other node in the head) device to provide node fencing. A STONITH device is a power switch that the cluster uses to reset nodes that are considered unresponsive. Resetting non-heartbeating nodes is the only reliable way to ensure that no data corruption is caused by nodes that hang and only appear to be dead. For more information about setting up STONITH, see, "Fencing and STONITH" (https://www.suse.com/documentation/sle_ha/book_sleha/data/book_sleha.html) in the *SUSE Linux Enterprise High Availability Administration Guide*.

Software Requirements for Creating a High Availability Environment

Ensure that the following software requirements are met for creating a high availability environment for the Cloud Manager Orchestration Server:

- SUSE Linux Enterprise Server 11 Service Pack 4 (SLES 11 SP 4) with all available online updates installed on each node that will be part of the cluster. Download the appropriate 64-bit version (https://download.suse.com/index.jsp?families=2658) from SUSE Downloads.
- SUSE Linux Enterprise High Availability Extension 11 Service Pack 4 (SLE HAE 11 SP 4) with all available online updates installed on each node that will be part of the cluster. Download the appropriate 64-bit version (https://download.suse.com/index.jsp?families=16189) from SUSE Downloads.

  The SLE HAE 11 SP 4 source includes Oracle Cluster File System 2 (OCFS2), a parallel cluster file system that offers concurrent access to a shared file system. See Section 15.3.4, "Setting Up the OCFS2 File System for High Availability," on page 141 for more information.

  The SLE HAE 11 SP 4 source also includes the Pacemaker Cluster Resource Manager. The Pacemaker software package is a high availability resource manager that supports multinode failover. See "Configuring Pacemaker" on page 140 for more information.

  SLE HAE 11 SP 4 integrates these open source storage technologies (OCFS2 and Pacemaker). This combined technology automatically shares cluster configuration and coordinates cluster-wide activities to ensure predictable administration of storage resources for shared-disk-based clusters.
- DNS installed on the nodes of the cluster for resolving the cluster host name to the cluster IP. See "The Domain Name System" (https://www.suse.com/documentation/sles11/book_sle_admin/data/cha_dns.html) in the *SUSE Linux Enterprise High Availability Administration Guide*.
- The Orchestration Server installed on each node of the cluster. See Section 7.2, "Configuring the Monitoring Server and Monitoring Agent," on page 62.

NOTE: Ensure that you install Cloud Manager Monitoring Server when you install Orchestration Server. Otherwise, later attempts to set up the server for high availability by running the `zos_server_ha_post_config.sh` script fail. See Section 3.3, “Cloud Manager Monitoring Server Installation Pattern,” on page 31.

Other Requirements

Ensure that the following services requirements are met for creating a high availability environment for the Orchestration Server:

- Synchronize time for nodes to an NTP server outside the cluster. See “Configuring Time Synchronization” on page 139.
- NIC names must be identical on all nodes.
- Use static IP addresses on each node and for the cluster.
- Include both the fully qualified host name and short host name in the `/etc/hosts` file.
- All cluster nodes must be able to access each other via SSH.

15.3.2 Installing the SLE High Availability Extension

After you install the SLES 11 SP 4 operating system, use YaST2 (or the command line, if you prefer) to install the SLE HAE 11 SP 4 package to each physical node that is to participate in the Orchestration Server cluster. See “Installation and Basic Setup” (https://www.suse.com/documentation/sle_ha/book_sleha/data/cha_ha_installation.html) in the *SUSE Linux Enterprise High Availability Administration Guide*.

15.3.3 Configuring Nodes with Time Synchronization and Installing Pacemaker to Each Node

After you have installed the high availability packages to each node of the cluster, you need to configure the Network Timing Protocol (NTP) and Pacemaker clustering environment on each physical machine that participates in the cluster.

- “Configuring Time Synchronization” on page 139
- “Configuring Pacemaker” on page 140

Configuring Time Synchronization

To configure time synchronization, you need to configure the nodes in the cluster to synchronize to a time server outside the cluster. The cluster nodes use the time server as their time synchronization source.

NTP is included as a network service in SLE HAE 11 SP 4. Use the *Time Synchronization with NTP* (https://www.suse.com/documentation/sles11/book_sle_admin/data/cha_netz_xntp.html) instructions in the *SUSE Linux Enterprise High Availability Administration Guide*. to help you configure each cluster node with NTP.

Configuring Pacemaker

Pacemaker Cluster Resource Manager is an open source server clustering system that ensures high availability and manageability of critical network resources including data, applications, and services. It is a multinode clustering product for Linux that supports failover, failback, and migration (load balancing) of individually managed cluster resources.

Pacemaker packages are installed with the SLE HAE 11 SP 4 install source. For detailed information about configuring Pacemaker, see *Configuring and Managing Cluster Resources (GUI)* (https://www.suse.com/documentation/sle_ha/book_sleha/data/cha_ha_configuration_gui.html) in the *SUSE Linux Enterprise High Availability Administration Guide*.

Setting timeout values is very important. The **Default Action Timeout** global cluster option controls how long Pacemaker waits for services to start if a specific timeout value is not configured otherwise in a cluster resource. The default value is 20 seconds.

Setting the value too low will result in unnecessary fencing operations in the cluster. Ensure that you check how long it takes Orchestration Server to start in your environment, and adjust the Default Action Timeout values accordingly.

NOTE: The Orchestration Server requires more time than 20 seconds to start. NetIQ Corporation recommends that you specify the value in this field to a minimum of `120s`. More time might be required if your Orchestration Server grid is very large.

To modify the Default Action Timeout value:

1. In the Pacemaker GUI, click **Configuration > CRM Config > Policy Engine**.
2. In the **Default Action Timeout** field, modify the default value of `20s` to `120s`.
3. Click **Apply**.

Figure 15-2 *The Main Settings Page in the Pacemaker Graphical Interface*

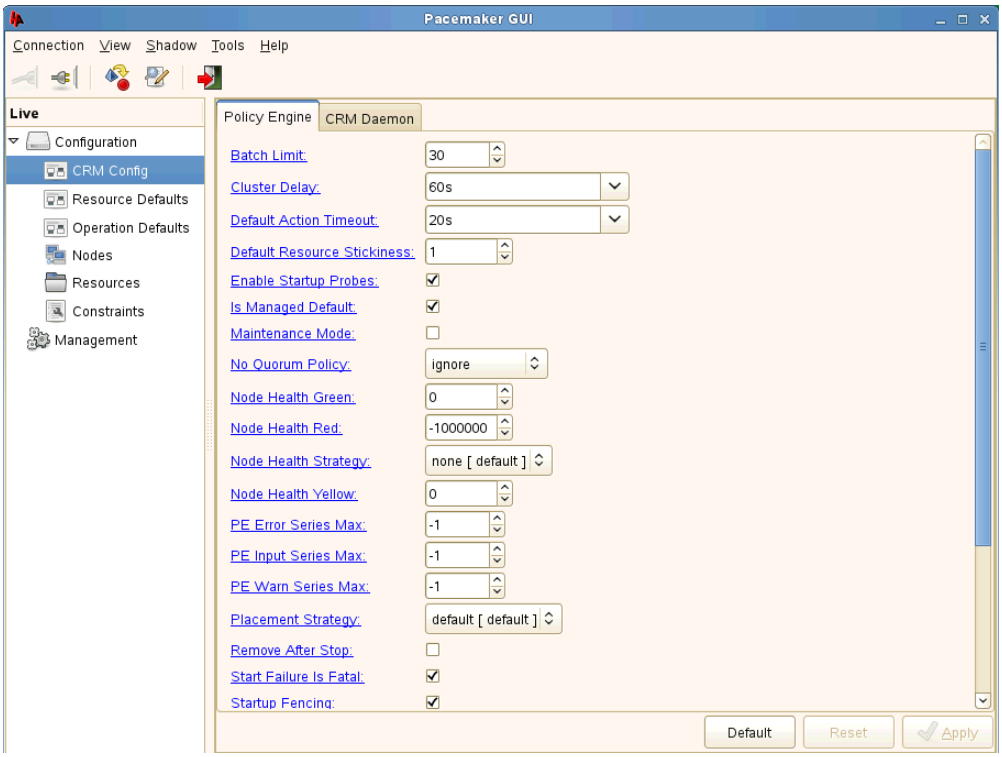

15.3.4 Setting Up the OCFS2 File System for High Availability

OCFS2 is a general-purpose journaling file system that is fully integrated in the Linux kernel that ships with SLE HAE. OCFS2 allows you to store application binary files, data files, and databases on devices on shared storage. All nodes in a cluster have concurrent read and write access to the file system. A distributed lock manager helps prevent file access conflicts. OCFS2 supports up to 32,000 subdirectories and millions of files in each directory. The O2CB cluster service (a driver) runs on each node to manage the cluster. OCFS2 uses the cluster membership services from the Pacemaker Cluster Resource Manager.

- "OCFS2 and Pacemaker" on page 141
- "Shared Storage Requirements for OCFS2" on page 141
- "DLM, O2CB, and STONITH" on page 141
- "Shared OCFS2 Volumes" on page 141

OCFS2 and Pacemaker

SUSE Linux Enterprise (SLE) High Availability Extension (HAE) server automatically installs the OCFS2 kernel module (`ocfs2`) and the Pacemaker Cluster Resource Manager. For each node, ensure that the `ocfs2-tools` and `ocfs2-kmp-*` packages are installed, and then configure the Pacemaker cluster management system.

Shared Storage Requirements for OCFS2

You must provide network storage where you will create the OCFS2 file system. This is where the Orchestration files will be stored. The High Availability Extension supports Fibre Channel or iSCSI storage area networks (SANs).

SAN configuration is beyond the scope of this document. Refer to the following resources:

- **Fibre Channel SAN:** See your vendor documentation.
- **iSCSI SAN:** See your vendor documentation. To configure an iSCSI SAN using SLES, see "Mass Storage over IP Networks: iSCSI (https://www.suse.com/documentation/sles11/stor_admin/data/cha_inst_system_iscsi.html)" in the *SLES Storage Administration Guide*.

DLM, O2CB, and STONITH

Before you can create shared OCFS2 volumes, you must configure the following resources as services in the cluster: DLM, O2CB, and a STONITH resource. For more information, see *"Configuring OCFS2 Services and a STONITH Resource"* (https://www.suse.com/documentation/sle_ha/book_sleha/data/sec_ha_ocfs2_create_service.html) in the *SLE High Availability Extension Administration Guide*.

Shared OCFS2 Volumes

For information about creating and using OCFS2 volumes, see *"OCFS2"* (https://www.suse.com/documentation/sle_ha/book_sleha/data/cha_ha_ocfs2.html) in the *SLE High Availability Extension Administration Guide*.

IMPORTANT: The Cloud Manager Orchestration Server requires a specific mount point for file storage on the SAN. Use `/zos` for this mount point.

15.3.5 Installing the Orchestration Server on the First Cluster Node

NOTE: As you prepare to install the Cloud Manager Orchestration Server and use it in a high availability environment, ensure that the requirements to do so are met. For more information, see Chapter 2, “Cloud Manager System Requirements,” on page 17.

To install the Orchestration Server packages on the first node of the cluster:

1. Log in to the target SLE HAE 11 SP 4 server as `root`, then open YaST2.
2. Download the appropriate Cloud Manager ISO to the node.

   or

   Load the Cloud Manager DVD on the node.
3. Define the Cloud Manager ISO or DVD as an add-on product:
   - **3a** In the YaST Control Center, click **Software,** then click **Add-On Products**.
   - **3b** Click Add, select **Local ISO Image** or **DVD**, then follow the prompts to add the product.
4. Read and accept the license agreement, then click **Next** to display the Software Selection and System Tasks dialog.

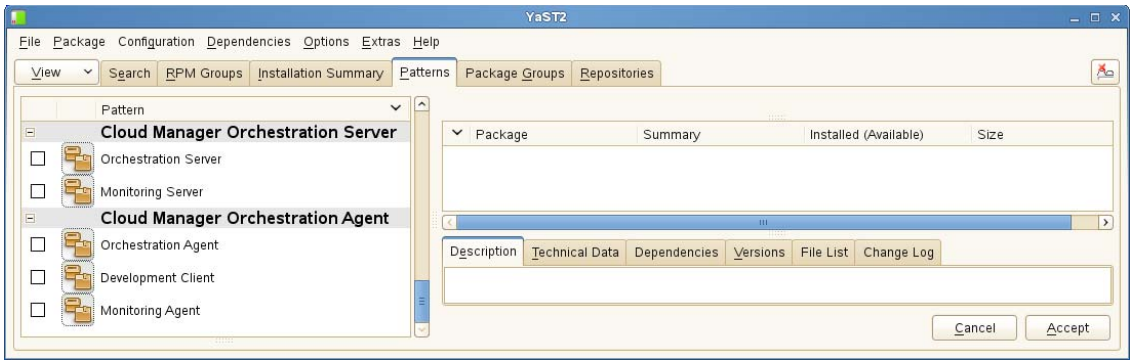

5. Select the Orchestration Server installation pattern for installation on the first node., then click **Accept**.

   When you select this pattern, the Monitoring Server installation pattern and the Monitoring Agent pattern are also selected. These patterns are the gateway between enterprise applications and resource servers. The Orchestration Server manages computing nodes (resources) and the jobs that are submitted from applications to run on these resources.

   TIP: If they are not already selected by default, you need to select the packages that are in the Orchestration Server pattern, the Monitoring Server pattern, and the Monitoring Client pattern.
6. Some additional packages that resolve the Orchestration Server dependencies are listed in an Automatic Changes dialog.

   Packages are written to your server.
7. When the package installation is complete, click **OK**.
8. Configure the Orchestration Server components that you have installed. You can use one of two methods to perform the configuration:
   - The Orchestration components (text-based) configuration script.
   - The Orchestration components GUI Configuration Wizard, which might be more user-friendly.

TIP: Although the text-based configuration process detects which RPM patterns are installed, the GUI Configuration Wizard requires that you specify which components are to be configured.

9 Finish the configuration by following the instructions in "Checking the Configuration" on page 146.

15.4 Configuring the Orchestration Server for High Availability

Configure the Orchestration Server that you installed on the first node of the cluster. Component configuration is done either with a text-based configuration tool or with a GUI Wizard configuration tool.

The text-based configuration script detects which RPM patterns are installed, but the GUI Configuration Wizard requires that you specify the components to be configured, whether the patterns have been installed on the server or not.

It is possible to execute the text-based configuration file Orchestration components from the Cloud Manager configuration utility, but this occurs only if you install Cloud Manager Application components on the same server as the Cloud Manager Orchestration components, which is only likely if you are setting up your system for a demonstration.

Both the text-based tool and the GUI Wizard tool produce a configuration file that can be used to automatically reconfigure your system after an upgrade. If you use the tools to reconfigure your server after the original configuration has been done, ensure that you reconfigure all of the components that are installed on the system (this is the default).

- Section 15.4.1, "Some Considerations When Configuring with the GUI Wizard," on page 144
- Section 15.4.2, "The Configuration Procedure," on page 144
- Section 15.4.3, "Checking the Configuration," on page 146
- Section 15.4.4, "Running the High Availability Configuration Script," on page 146

When you have configured the Orchestration Server, you need to complete the other items necessary for a high availability setup in the following order:

1. Section 15.5, "Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster," on page 147.
2. Section 15.6, "Creating the Server Cluster Resource Group," on page 147.
3. Section 15.7, "Testing the Failover of the Orchestration Server in a High Availability Grid," on page 148
4. Section 15.8, "Installing and Configuring other Orchestration Components to the High Availability Grid," on page 148.

15.4.1 Some Considerations When Configuring with the GUI Wizard

If you have only a keyboard to navigate through the pages of the GUI Configuration Wizard, use the Tab key to shift the focus to a control you want to use (for example, a **Next** button), then press the Spacebar to activate this control.

When you have finished answering the configuration questions in the wizard, the Cloud Manager Orchestration Configuration Summary page displays. Although this page of the wizard lets you navigate by using the Tab key and the Spacebar, you need to use the Ctrl+Tab combination to navigate past the summary list. Click **Back** if you accidentally enter the summary list, and re-enter the page to navigate to the control buttons.

By default, the **Configure now** check box on the page is selected. If you accept this default, the wizard starts the Orchestration Server and applies the configuration settings. If you deselect the check box, the wizard writes out the configuration file to `/etc/opt/novell/novell_zenworks_orch_install.conf` without starting the Orchestration Server or applying the configuration settings.

You can use this `.conf` file to start the Orchestration Server or Agent and apply the settings either manually or with an installation script. Use the following command to run the configuration:

```
/opt/novell/zenworks/orch/bin/config -rs
```

When the installation and configuration are complete, you need to validate and optimize the configuration.

15.4.2 The Configuration Procedure

To configure the Orchestration Server for use in a high-availability environment,

1. Ensure that you are logged in as `root` to run the configuration.
2. Ensure that you are ready with the information that you will be prompted for during the configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Configuration Type	Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation. This section discusses standard installation, so specify `h` (for `ha` which means “high availability”).
Cluster Hostname or IP Address	Specify the fully qualified cluster host name or the IP address that is used for configuring the Orchestration Server instance in a high availability cluster. The configuration script binds the IP address of the cluster to this server.
Grid Name	A grid is an administrative domain container holding all of the objects in your network or data center. The Orchestration Server monitors and manages these objects, including users, resources, and jobs. The grid name you create here is displayed as the name for the container placed at the root of the Explorer tree in the Orchestration Console.

Server Configuration Requirement	Explanation and Action
Administrator User	Specify a name for the Orchestration Server Administrator user. This name is used to log in as the administrator of the Orchestration Server and the objects it manages.
Administrator Password	Specify a password for the Orchestration Administrator user, then retype the password to validate it. You should remember this user name for future logins.
Path to License File	A license key (90-day evaluation license or a full license) is required to use this product. You should have received this key from NetIQ, then you should have subsequently copied it to the network location that you specify here. Be sure to include the name of the license file in the path.
Auditing Database	We recommend that you do not install the audit database on this server.
Orchestration Agent Port[1]	Port 8100 is used for communication between the Orchestration Server and the Orchestration Agent. Specify another port number if 8100 is reserved for another use. If your Orchestration Server communicates with ESX servers, we recommend you configure port 8101. This requires that you configure all other Orchestration Agents communicating with this server to use port 8101. This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the **Configure Advanced Settings** check box in the wizard, you have the option of changing the default values. If you leave the check box deselected the setting is configured with normal defaults.
Administrator Information Port[1]	Port 8001 on the Orchestration Server provides access to an Administrator Information page that includes links to product documentation, agent and client installers, and product tools to help you understand and use the product. Specify another port number if 8001 is reserved for another use on this server.
TLS Certificate and Key[1]	Choose whether to generate a TLS certificate and key. ◆ Default = `yes` (the Orchestration Server must generate a certificate and key for authentication) ◆ A PEM-encoded TLS certificate and key is needed for secure communication between the Orchestration Server and Orchestration Agent. ◆ If you respond with `no`, you need to provide the location of an existing certificate and key.
TLS Server Certificate[2]	Specify the full path to the TLS server certificate. ◆ Default = `/etc/ssl/servercerts/servercert.pem` ◆ Specify the path to the existing TLS certificate.
TLS Server Key[2]	Specify the full path to the TLS server private key. ◆ Default = `/etc/ssl/servercerts/serverkey.pem` ◆ Specify the path to the existing TLS private key.

[1] This configuration parameter is considered an advanced setting for the Orchestration Server in the Orchestration Components Configuration Wizard. If you select the **Configure advanced settings** check box in the wizard, the setting is configured with normal defaults. Leaving the check box deselected lets you have the option of changing the default value.

[2] This configuration parameter is considered an advanced setting for the Orchestration Server in the Orchestration Components Configuration Wizard. If you select the **Configure advanced settings** check box in the wizard, this parameter is listed, but default values are provided only if the previous value is manually set to no.

3 At the computer where you installed the Cloud Manager Orchestration Server pattern, run the Cloud Manager Orchestration configuration utility:

```
/opt/novell/zenworks/orch/bin/config
```

or

```
/opt/novell/zenworks/orch/bin/guiconfig
```

4 Continue with "Checking the Configuration" on page 146.

15.4.3 Checking the Configuration

When the configuration is completed, the first node of the Orchestration Server cluster is set up. You then need to check the configuration.

1 Open the configuration log file (`/var/opt/novell/novell_zenworks_orch_install.log`) to ensure that the components were correctly configured.

You can change the configuration if you change your mind about some of the parameters you provided in the configuration process. To do so, rerun the configuration and change your responses.

The configuration tool performs the following functions in sequence on the Orchestration Server:

1. Binds the cluster IP on this server by issuing the following command internally:

```
IPaddr2 start <IP_address_you_provided>
```

IMPORTANT: Ensure that you configure DNS to resolve the cluster host name to the cluster IP.

---

2. Configures the Orchestration Server.
3. Shuts down the Orchestration Server because you specified that this is a high availability configuration
4. Unbinds the cluster IP on this server by issuing the following command internally:

```
IPaddr2 stop <IP_address_you_provided>
```

2 Continue with "Running the High Availability Configuration Script" on page 146.

15.4.4 Running the High Availability Configuration Script

Before you run the high availability configuration script, ensure that you have installed the Orchestration Server to a single node of your high availability cluster. For more information, see Section 15.3.5, "Installing the Orchestration Server on the First Cluster Node," on page 142

IMPORTANT: The high availability configuration script asks for the mount point on the Fibre Channel SAN. Ensure that you have that information (`/zos`) before you run the script.

The high availability script, `zos_server_ha_post_config.sh`, is located in `/opt/novell/zenworks/orch/bin/ha` with the other configuration tools. You need to run this script on the first node of the cluster (that is, the node where you installed the Orchestration Server) as the next step in setting up Cloud Manager Orchestration Server to work in a high availability environment.

The script performs the following functions:

- Verifies that the Orchestration Server is not running
- Copies Apache files to shared storage
- Copies gmond and gmetad files to shared storage
- Moves the Orchestration files to shared storage (first node of the cluster)
- Creates symbolic links pointing to the location of shared storage (all nodes of the cluster)

The high availability configuration script must be run on all nodes of the cluster. Ensure that you follow the prompts in the script exactly; do not misidentify a secondary node in the cluster as the primary node.

15.5 Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster

After you have followed the steps to set up the primary node in your planned cluster, you need to set up the other nodes that you intend to use for failover in that cluster. Use the following sequence as you set up other cluster nodes (the sequence is nearly identical to setting up the primary node):

15.6 Creating the Server Cluster Resource Group

The resource group creation script, `zos_server_ha_resource_group`, is located in `/opt/novell/zenworks/orch/bin/ha` with the other configuration tools. You can run this script on the first node of the cluster to set up the cluster resource group.

The script performs the following functions:

- Obtains the DNS name from the Orchestration Server configuration file.
- Creates the cluster resource group.
- Configures resource stickiness to avoid unnecessary failback operations.

When you have installed and configured the nodes in the SLES 11 SP 4 cluster and created a cluster resource group, use the Pacemaker tools to start the cluster resource group. For more information, see “Cluster Management Tools (http://www.suse.com/documentation/sle_ha/book_sleha/data/cha_ha_management.html)” in the *SUSE Linux Enterprise High Availability Extension Guide*.

You are then ready to test the failover of the Orchestration Server in the high-availability cluster (see Section 15.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 148).

15.7 Testing the Failover of the Orchestration Server in a High Availability Grid

You can optionally simulate a failure of the Orchestration Server by powering off or performing a shutdown of the server. After approximately 30 seconds, the clustering software detects that the primary node is no longer functioning, binds the IP address to the failover server, then starts the failover server in the cluster.

Access the Orchestration Administrator Information Page to verify that the Orchestration Server is installed and running (stopped or started). Use the following URL to open the page in a web browser:

`http://DNS_name_or_IP_address_of_cluster:8001`

The Administrator Information page includes links to separate installation programs (installers) for the Orchestration Agent and the Orchestration Clients. The installers are used for various operating systems.

15.8 Installing and Configuring other Orchestration Components to the High Availability Grid

To install and configure other Orchestration components (including the Orchestration Agent, the Monitoring Agent, or the Monitoring Server) on servers that authenticate to the cluster, you need to determine which components you want to install, remembering these dependencies:

- Orchestration components must be installed on platforms that are tested and supported. For more information, see Chapter 2, “Cloud Manager System Requirements,” on page 17.
- Use YaST2 to install the Orchestration packages of your choice to the network server resources of your choice. For more information, see Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 39.

  If you want to, you can download the Orchestration Agent or clients from the Administrator Information page and install them to a network resource.
- Run the text-based configuration script or the GUI Configuration Wizard to configure the Orchestration components you have installed (including any type of installation of the agent). As you do this, you need to remember the host name of the Orchestration Server (that is, the primary Orchestration Server node), and the administrator name and password of this server. For more information, see Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 39.

It is important to understand that virtual machines under the management of the Cloud Manager Orchestration Server are also highly available—the loss of a host causes the Orchestration Server to re-provision it elsewhere. This is true as long as the constraints in the Orchestration Server allow it to re-provision (for example, if the virtual machine image is on shared storage).

16 Configuring the Orchestration Server to Use an Audit Database

When you install Cloud Manager Orchestration Server, you can optionally point it to a relational database that you can use to audit the work done by the product. There is no relational database management system bundled with the product, but because Orchestration Server runs on SLES 11 SP 4, you can use a PostgreSQL database that ships with SLES and configure it for use with Orchestration Server auditing. If you want to use another database, you must configure it separately for use with the Orchestration Server.

- Section 16.1, “Installing the PostgreSQL Package and Dependencies on an Independent Host,” on page 149
- Section 16.2, “Configuring PostgreSQL to Accept Remote Database Connections,” on page 151
- Section 16.3, “Logging in Locally to the PostgreSQL Database,” on page 152
- Section 16.4, “Creating an Orchestration Server User for the PostgreSQL Database,” on page 152
- Section 16.5, “Configuring the Orchestration Server Audit Database on a Separate Host,” on page 153
- Section 16.6, “Installing and Configuring the Orchestration Server for Use with a Local PostgreSQL Audit Database,” on page 154
- Section 16.7, “Configuring the Audit Database after the Cloud Manager Orchestration Server Is Configured,” on page 157
- Section 16.8, “Configuring the Remote Audit Database after the Cloud Manager Orchestration Server Is Configured,” on page 158
- Section 16.9, “Modifying Audit Database Tables to Accommodate Long Names,” on page 159
- Section 16.10, “Understanding Grid ID Usage in the Audit Database,” on page 159

16.1 Installing the PostgreSQL Package and Dependencies on an Independent Host

When you enable and configure Orchestration Server auditing, you create a small custom database and a simple schema that persists all of the Orchestration Server jobs that have been run, along with their parameters. The database also maintains the login or logout activity of the Orchestration Server users and resources and includes an “actions” table that records provisioning actions and their status (started, failed, completed successfully, etc.).

NOTE: We recommend that you install the PostgreSQL packages on a different server than the one where you install the Orchestration Server. This ensures an adequate amount of space for running the server as the database is used.

We also recommend that you open TCP port 5432 (or whatever port you configure PostgreSQL to use—5432 is the PostgreSQL default) in the firewall of the RDBMS host. Without an open port in the host firewall, a remote Orchestration Server cannot access the audit database.

For high availability Orchestration Server configurations, you need to install the database outside of the high availability cluster.

If you want to run the database on the same host with the Orchestration Server, see Section 16.6, “Installing and Configuring the Orchestration Server for Use with a Local PostgreSQL Audit Database,” on page 154.

If the Orchestration Server does not have PostgreSQL packages installed and running and you are not using an external PostgreSQL database instance, use YaST to search for `postgresql-server`, then install the package and its dependencies.

You can also run the following command from the bash prompt:

```
yast2 -i postgresql-server
```

When PostgreSQL is installed, you need to create the default database and start it. Use the following commands:

```
su - postgres
```

```
initdb
```

```
pg_ctl start
```

These commands create or update the PostgreSQL privilege database and installs the prepared tables. For more detail about what you will see when you run these commands, see “Detail” on page 150.

NOTE: You cannot run the `pg_ctl` command as `root`. You must first change to the superuser for PostgreSQL (`su - postgres`). Failure to issue this command first results in the following messages:

```
# pg_ctl start
pg_ctl: cannot be run as root
Please log in (using, e.g., "su") as the (unprivileged) user that will
own the server process.
```

16.1.1 Detail

```
postgres> initdb
The files belonging to this database system will be owned by user "postgres".
This user must also own the server process.

The database cluster will be initialized with locale en_US.UTF-8.
The default database encoding has accordingly been set to UTF8.

creating directory /var/lib/pgsql/data ... ok
creating directory /var/lib/pgsql/data/global ... ok
creating directory /var/lib/pgsql/data/pg_xlog ... ok
creating directory /var/lib/pgsql/data/pg_xlog/archive_status ... ok
creating directory /var/lib/pgsql/data/pg_clog ... ok
creating directory /var/lib/pgsql/data/pg_subtrans ... ok
creating directory /var/lib/pgsql/data/pg_twophase ... ok
creating directory /var/lib/pgsql/data/pg_multixact/members ... ok
creating directory /var/lib/pgsql/data/pg_multixact/offsets ... ok
creating directory /var/lib/pgsql/data/base ... ok
creating directory /var/lib/pgsql/data/base/1 ... ok
creating directory /var/lib/pgsql/data/pg_tblspc ... ok
selecting default max_connections ... 100
```

```
selecting default shared_buffers ... 1000
creating configuration files ... ok
creating template1 database in /var/lib/pgsql/data/base/1 ... ok
initializing pg_authid ... ok
enabling unlimited row size for system tables ... ok
initializing dependencies ... ok
creating system views ... ok
loading pg_description ... ok
creating conversions ... ok
setting privileges on built-in objects ... ok
creating information schema ... ok
vacuuming database template1 ... ok
copying template1 to template0 ... ok
copying template1 to postgres ... ok

WARNING: Enabling "trust" authentication for local connections
You can change this by editing pg_hba.conf or using the -A option the
next time you run initdb.

Success. You can now start the database server using:

    postmaster -D /var/lib/pgsql/data
or
    pg_ctl -D /var/lib/pgsql/data -l logfile start

postgres> postmaster -i
```

16.2 Configuring PostgreSQL to Accept Remote Database Connections

To configure the PostgreSQL database to accept remote database connections, you need to add the following line to the `/var/lib/pgsql/data/pg_hba.conf` file:

```
host    all      all         0.0.0.0/0      trust
```

NOTE: After initial configuration, you can replace the `0.0.0.0/0` with a more restrictive mask. In a high availability server configuration, ensure that each host in the high availability cluster is enabled as a remote host.

For added security, the `/var/lib/pgsql/data/pg_hba.conf` file should list only the desired hosts. For example, only the Orchestration Server would be included in the `trust` line.

After you make the change to the `pg_hba.conf` file, you need to specify the following command so that you do not receive an error when remote hosts try to connect:

```
pg_ctl reload
```

If `pg_hba.conf` is not configured and you attempt to connect, an error similar to the following is displayed:

```
psql: FATAL:  no pg_hba.conf entry for host "10.99.15.64", user "postgres",
database "postgres", SSL off
```

Depending on the environment, you might need to perform some additional configuration for remote database setup. Editing the `listen_addresses` section of the `postgresql.conf` file enables the database server to listen for incoming connections on the specified IP addresses. The following is an excerpt from that section of the file:

```
listen_addresses = 'localhost'
                   # what IP address(es) to listen on;
                   # comma-separated list of addresses;
                   # defaults to 'localhost', '*' = all
```

After you modify the `listen_addresses` entry in `postgresql.conf`, use the following command to restart the PostgreSQL server (recommended in the PostgreSQL documentation):

```
pg_ctl restart
```

16.3 Logging in Locally to the PostgreSQL Database

When you have installed the database, the next step is to check that you can connect to the database on the database host. The default admin user name is `postgres`. Use the following commands to set up a password for the postgres user on the database host machine:

```
psql
```

NOTE: Remember the password. You need to use it later to log in to the database.

Running this command results in a screen like this:

```
Welcome to psql 8.1.11, the PostgreSQL interactive terminal.

Type:  \copyright for distribution terms
       \h for help with SQL commands
       \? for help with psql commands
       \g or terminate with semicolon to execute query
       \q to quit

postgres=# alter user postgres password 'pass';
ALTER ROLE
postgres=#
```

16.4 Creating an Orchestration Server User for the PostgreSQL Database

Next, set up a PostgreSQL user to own the audit database schema before you run the server configuration script or the GUI Configuration Wizard.

1. On the database host machine, use the following commands to log in as `root` at the database host machine:

   ```
   su - postgres
   ```

   ```
   psql
   ```

2. At the `psql` prompt on the database host, use the following command to create an audit database schema user, for example:

   ```
   postgres=# create user zos password 'zos';
   ```

   ```
   CREATE ROLE
   ```

   Single quotes surrounding the password are required.

3. Enter the `\q` command at the `psql` prompt to exit the database.

16.5 Configuring the Orchestration Server Audit Database on a Separate Host

The easiest way to configure the audit database is to do so when you configure the Orchestration Server. Use the following procedure to configure the database.

NOTE: The questions presented in the text-based config script are shown here, but the questions presented in the graphical Configuration Wizard are similar.

---

1. After you have installed the Cloud Manager Orchestration packages you want, run the configuration (either the config script or the graphical Configuration Wizard) until you see the following question:

   ```
   Enable Auditing (y/n) [no]:
   ```

2. Enter `yes` to answer this question. The following question displays:

   ```
   Configure Audit DB (y/n) [no]:
   ```

3. Enter `yes` to answer this question. The following question displays:

   ```
   Jdbc URL [jdbc:postgresql://localhost/]:
   ```

4. Enter the URL of the server where PostgreSQL is running, then press Enter.

   `jdbc:postgresql://`*`IP_address_of_database_server/`*

   This is a standard JDBC URL because this is a Java server that uses JDBC for the interface database. The URL must be properly formed, with a slash and without a database name at the end. We do not recommend using “localhost” as the URL.

   The following prompt is displayed:

   ```
   DB Admin Username:
   ```

5. Specify the PostgreSQL database administrator user name, then press Enter.

   This is the same user name that was created when PostgreSQL was installed. In most instances, the user name is `postgres`.

   The following prompt is displayed:

   ```
   DB Admin Password:
   ```

6. Specify the PostgreSQL database administrator password, then press Enter.

   The following prompt is displayed:

   ```
   Retype password:
   ```

7. Retype the database administrator password to verify it, then press Enter. The following prompt is displayed:

   ```
   ZOS Audit Database Name [zos_db]:
   ```

8. Specify the name of the database you want to create for Orchestration Server auditing, then press Enter. The following prompt is displayed:

   ```
   Audit DB Username:
   ```

9 Specify the name you want to use for the PostgreSQL database user that will be used by the Orchestration Server for auditing (that is, a user with Read and Write privileges, not the administrator), then press Enter. The following prompt is displayed:

```
Audit DB Password:
```

10 Specify the password you want to use for authentication by the designated PostgreSQL database user, then press Enter. The following prompt is displayed:

```
Retype password:
```

11 Retype the password, then press Enter.

After you retype the new audit database password, the configuration interview for the Orchestration Server continues normally.

16.6 Installing and Configuring the Orchestration Server for Use with a Local PostgreSQL Audit Database

When you install the Cloud Manager Orchestration Server, you can optionally point it to a relational database that you can use to audit the work done by the product. There is no relational database management system bundled with the product, but because the Orchestration Server is supported on a SLES 11 SP 4 server, you can use a PostgreSQL database that is delivered as part of the server operating system, and configure it for use with Orchestration Server auditing. If you want to use some other database, you must configure it separately for use with Cloud Manager.

- Section 16.6.1, “Installing the PostgreSQL Package and Dependencies,” on page 154
- Section 16.6.2, “Configuring PostgreSQL to Accept Local Database Connections,” on page 155
- Section 16.6.3, “Logging in Locally to the PostgreSQL Database,” on page 155
- Section 16.6.4, “Installing and Configuring the Local Orchestration Server Audit Database,” on page 156

16.6.1 Installing the PostgreSQL Package and Dependencies

NOTE: We recommend that you install the PostgreSQL package separately on different server from the one where you install the Cloud Manager Orchestration Server. Using a separate server helps ensure an adequate amount of space for running the server as the database is used.

For more information, see Section 16.2, “Configuring PostgreSQL to Accept Remote Database Connections,” on page 151.

If your server does not have the PostgreSQL package installed and running and you do not want to use a separate server for the database, use YaST to search for `postgresql-server`, then you can install the package and its dependencies on the Orchestration Server.

You can also run the following command from the bash prompt:

```
yast2 -i postgresql-server
```

When PostgreSQL is installed, you need to create the default database and start it. Use the following commands:

```
su - postgres
```

```
initdb

pg_ctl start
```

These commands create or update the PostgreSQL privilege database and install the prepared tables. For more detail about what you will see when you run these commands, see “Detail” on page 150.

NOTE: You cannot run the `pg_ctl` command as `root`. You must first change to the superuser for PostgreSQL (`su - postgres`). Failure to issue this command first results in the following messages:

```
# pg_ctl start
pg_ctl: cannot be run as root
Please log in (using, e.g., "su") as the (unprivileged) user that will
own the server process.
```

16.6.2 Configuring PostgreSQL to Accept Local Database Connections

To configure the PostgreSQL database to accept remote database connections, you need to change the following line in the `/var/lib/pgsql/data/pg_hba.conf` file:

```
host    all      all         0.0.0.0/0      ident sameuser
```

The line should be changed as follows:

```
host    all      all         0.0.0.0/0      trust
```

16.6.3 Logging in Locally to the PostgreSQL Database

When you have installed the database, the next step is to check that you can connect to the database on the database host. The default admin user name is `postgres`. Use the following commands to set up a password for the postgres user on the database host machine:

```
psql
```

NOTE: Remember the password. You need it to log in to the database later.

Running this command results in a screen similar to this:

```
Welcome to psql x.x.x, the PostgreSQL interactive terminal.

Type:  \copyright for distribution terms
       \h for help with SQL commands
       \? for help with psql commands
       \g or terminate with semicolon to execute query
       \q to quit

postgres=# alter user postgres password 'pass';
ALTER ROLE
postgres=#
```

16.6.4 Installing and Configuring the Local Orchestration Server Audit Database

When you enable and configure Orchestration Server auditing, you create a small custom database and a simple schema that persists all of the Orchestration Server jobs that have been run, along with their parameters.The database also maintains the login or logout activity of the Cloud Manager users and resources.

The easiest way to configure the audit database is to do so when you configure the Orchestration Server. Use the following procedure to configure the database.

NOTE: The questions presented in the text-based config script are shown here, but the questions presented in the graphical Configuration Wizard are similar.

---

1. After you have installed the Cloud Manager packages you want, run the configuration (either the config script or the graphical Configuration Wizard) until you see the following question:

   ```
   Enable Auditing (y/n) [no]:
   ```

2. Enter `yes` to answer this question. The following question displays:

   ```
   Configure Audit DB (y/n) [no]:
   ```

3. Enter `yes` to answer this question. The following question displays:

   ```
   Jdbc URL [jdbc:postgresql://localhost/]:
   ```

4. Press Enter to accept the default (`jdbc:postgresql://localhost/`).

   This is a standard JDBC URL because this is a Java server that uses JDBC for the interface database. The URL must be properly formed, with a slash and without a database name at the end.

   The following prompt is displayed:

   ```
   DB Admin Username:
   ```

5. Specify the PostgreSQL database administrator user name, then press Enter.

   This is the same name that was specified when PostgreSQL was installed. In most instances, the user name is `postgres`.

   The following prompt is displayed:

   ```
   DB Admin Password:
   ```

6. Specify the PostgreSQL database administrator password, then press Enter.

   The following prompt is displayed:

   ```
   Retype password:
   ```

7. Retype the database administrator password to verify it, then press Enter. The following prompt is displayed:

   ```
   ZOS Audit Database Name [zos_db]:
   ```

8. Specify the name of the database you want to create for Orchestration Server auditing, then press Enter. The following prompt is displayed:

   ```
   Audit DB Username:
   ```

9 Specify the name you want to use for the PostgreSQL database user that will be used by the Orchestration Server for auditing (that is, a user with Read and Write privileges, not the administrator), then press Enter. The following prompt is displayed:

```
Audit DB Password:
```

10 Specify the password you want to use for authentication by the designated PostgreSQL database user, then press Enter. The following prompt is displayed:

```
Retype password:
```

11 Retype the password, then press Enter.

After you retype the new audit database password, the configuration interview for the Orchestration Server continues normally.

16.7 Configuring the Audit Database after the Cloud Manager Orchestration Server Is Configured

If you have already installed and configured the Cloud Manager Orchestration Server, it is still possible to configure an audit database.

1 On the Orchestration Server host machine, use your favorite editor to edit the script `/opt/novell/zenworks/zos/server/conf/audit_db_prep.sql`:

1a Replace the `${DB_NAME}` variable with the PostgreSQL database name (for example, `zos_db`).

1b Replace the `${DB_USER}` variable with the PostgreSQL schema owner name (for example, `zos`).

2 Use the following commands to run the modified script as the PostgreSQL database administrator:

```
su - postgres
psql -f audit_db_prep.sql
```

3 Use the following command to log into PostgreSQL, using the database name and schema owner substituted in Step 1 above:

```
su - postgres
psql -d zos_db -U zos -f audit_db_def.sql
```

4 Confirm that the database user name and password match the values used when creating the schema owner database user in Section 16.4, “Creating an Orchestration Server User for the PostgreSQL Database,” on page 152. In this example, the user name is `zos` and the password is `zos`.

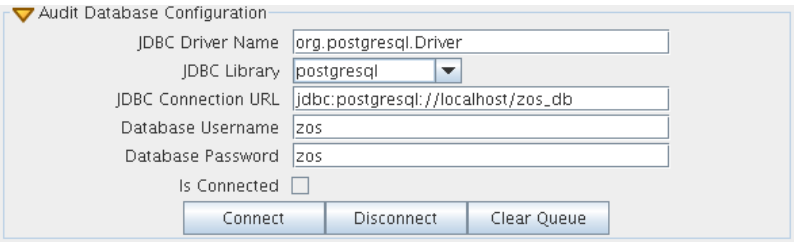

5 Confirm that the database user name and password match the values you replaced in the variables of the `.sql` script. In this example, the user name is `zos` and the password is `zos`.

6 Click **Connect**.

The **Is Connected** check box is selected; the Orchestration Server is connected to the database so that any queued data and subsequent job, user, and resource events are written there.

16.8 Configuring the Remote Audit Database after the Cloud Manager Orchestration Server Is Configured

If you have already installed and configured the Cloud Manager Orchestration Server, it is still possible to configure an audit database.

1 On the Orchestration Server host machine, use your favorite editor to edit the script `/opt/novell/zenworks/zos/server/conf/audit_db_def.sql`:

1a Replace the `${DB_NAME}` variable with the PostgreSQL database name (for example, `zos_db`).

1b Replace the `${DB_USER}` variable with the PostgreSQL schema owner name (for example, `zos`).

2 Use the following commands to run the modified script as the PostgreSQL database administrator for the remote database:

```
su - postgres
psql -h <psql-server-addr> -d postgres -U postgres -f audit_db_prep.sql
```

3 Use the following command to log into PostgreSQL, using the database name and schema owner substituted in Step 1 above:

```
su - postgres
psql -h <psql-server-addr> -d zos_db -U zos -f audit_db_def.sql
```

4 Confirm that the database user name and password match the values used when creating the schema owner database user in Section 16.4, “Creating an Orchestration Server User for the PostgreSQL Database,” on page 152. In this example, the user name is `zos` and the password is `zos`.

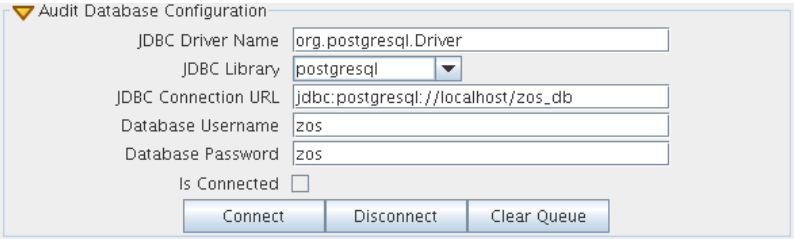

5 Confirm that the database user name and password match the values you replaced in the variables of the `.sql` script. In this example, the user name is `zos` and the password is `zos`.

6 Click **Connect**.

The **Is Connected** check box is selected; the Orchestration Server is connected to the database so that any queued data and subsequent job, user, and resource events are written there.

16.9 Modifying Audit Database Tables to Accommodate Long Names

If your installation of the Cloud Manager Orchestration Server uses Grid Object names that have an unusual number of characters, the server might lose its connection with the audit database.

If your Grid Objects are named with long names, you might need to configure some of the table columns in the audit database with different sizes. Here are some things you need to know about the database and how to make such changes:

- The default length of some names is predefined in the audit database. For example, the `user name` and the `resource name` size in the audit database both default to 30 characters in allowable length.
- The workflow ID (`originWorkflowId`, `parentWorkflowId`) in the workflow table is constructed by concatenating the name of the user who invoked the job + the name of the deployed job + an instance number. The default size value is 100.
- The job instance ID (`jobinstanceid`) in the workflow table includes either the deployed name of the job or a server component name that invoked the job. For example, when the Scheduler invokes the job, then `Scheduler` is concatenated with the deployed job name. For example: `Scheduler(cpuinfo)`.
- The `name` column in the sessions table records both user and resource names.

The SQL table definition is found at `<server>/conf/audit_db_def.sql`.

Use SQL commands to change an existing table column. The following excerpts from the database show some table columns that you might need to change:

```
CREATE TABLE    actions          (

                targetobjectname        VARCHAR(50)      NOT NULL,
                username                VARCHAR(30)      NOT NULL,
                jobinstanceid           VARCHAR(100)


CREATE TABLE workflow (

       jobId VARCHAR(100) NOT NULL,
       jobInstanceName VARCHAR(100) NOT NULL,
       deployedJobName VARCHAR(30) NOT NULL,
       originWorkflowId VARCHAR(100) NOT NULL,
       parentWorkflowId VARCHAR(100),
       username VARCHAR(30) NOT NULL,



CREATE TABLE sessions (

       name VARCHAR(30) NOT NULL,
```

16.10 Understanding Grid ID Usage in the Audit Database

The Orchestration Grid ID is created using either the `config` or `guiconfig` configuration wizard operations or the `zosadmin create -g` command. The grid name you specify is displayed as the name for the container placed at the root of the tree in the Explorer panel of the Orchestration

Console.The value for the Grid ID is saved in the `/var/opt/novell/zenworks/zos/server/zos.conf` file by the property `system.property.com.novell.zos.server.gridId`. Historical records of job instances (also known as “workflows”) run on the Orchestration grid are stored in the audit database (if included in the Orchestration Server installation) and are indexed by Grid ID. If you change the value of the Grid ID, the Orchestration Console loses access to these records.

For instance, testing has shown that if you choose to upgrade the Orchestration Server using the `zosadmin create --upgrade -g` option (that is selecting a Grid ID) instead of the config or guiconfig operations, it is possible that you might not use the existing Grid ID value or that you might neglect to use a value with the command. In this case, the default (the fully-qualified domain name of the current host) is used, which could differ from the original value for the Grid ID.

If this happens, any workflows recorded in the audit database prior to the upgrade are not displayed in the Orchestration Console, but they are still recorded in the **gridid** column of the workflows table in the database.

NOTE: You can use an SQL query if you want to retrieve the workflows. The first part of such a query might look like this:

```
SELECT * FROM workflow WHERE gridId = 'labzos.pso.lab.novell.com_Grid' AND ...
```

Keep in mind that the original Grid ID is not lost. If you want a report on that ID, you can use the `zosadmin audit*` commands with a `-g` option to yield a report showing the old Grid ID.

If you want to change the Grid ID, you have several options:

- Edit `/var/opt/novell/zenworks/zos/server/zos.conf` and change the value of the Grid ID.
- Change the ID during an upgrade using `zosadmin create --upgrade -g` and with a new value for the Grid ID.
- Use SQL commands to change the Grid ID of existing records in the audit database. For example,

  ```
  UPDATE actions SET gridid = 'newgridname' WHERE gridid = 'oldname';
  UPDATE workflow SET gridid = 'newgridname' WHERE gridid = 'oldname';
  UPDATE sessions SET gridid = 'newgridname' WHERE gridid = 'oldname';
  ```

17 Integrating the Orchestration Server with a Sentinel Collector

This section provides details about integrating the NetIQ Sentinel collector for the Cloud Manager Orchestration Server with your Cloud Manager and NetIQ Sentinel installations to make the Orchestration Server logging easier to view and interpret.

This integration is not required for the Orchestration Server to function. Logging with NetIQ Sentinel is entirely independent of the audit database feature. For more information about installing the audit database feature, see Section 16.1, “Installing the PostgreSQL Package and Dependencies on an Independent Host,” on page 149.

- Section 17.1, “Integration Architecture,” on page 161
- Section 17.2, “System Requirements,” on page 162
- Section 17.3, “Importing and Deploying the Orchestration Server Sentinel Collector Plug-in,” on page 163
- Section 17.4, “Connecting the Orchestration Server to the Sentinel Collector Plug-In,” on page 165
- Section 17.5, “Verifying the Sentinel Configuration After Connecting to the Orchestration Server,” on page 165
- Section 17.6, “Event Classification and Taxonomy Keys,” on page 166
- Section 17.7, “Plain Text Visibility of Sensitive Information,” on page 169

17.1 Integration Architecture

NetIQ Sentinel is a security information and event management solution that receives information from many sources throughout an enterprise, then standardizes the information, prioritizes it, and presents it to you so that you can make threat, risk, and policy-related decisions. The Sentinel Control Center is the main user interface for viewing and interpreting this data. For overall information about NetIQ Sentinel, see the NetIQ Sentinel Documentation website (http://www.novell.com/documentation/sentinel).

The Orchestration Server can be configured to send log events to Sentinel over a single SSL connection (typically port 1443). The events are sent in RFC5424 (syslog) format, and are received by the Sentinel Event Source Server, which, for each event, parses the syslog header, and then hands the event over to the Orchestration Server Collector plug-in for Sentinel. The Sentinel collector parses the encapsulated Orchestration Server log event and performs normalization tasks before finally submitting it to the Sentinel event processing engine. These normalization tasks include mapping Orchestration Server log levels to Sentinel numerical event severities and extracting event metadata.

Figure 17-1 *Simplified Architecture for Orchestration Server Collector Integration*

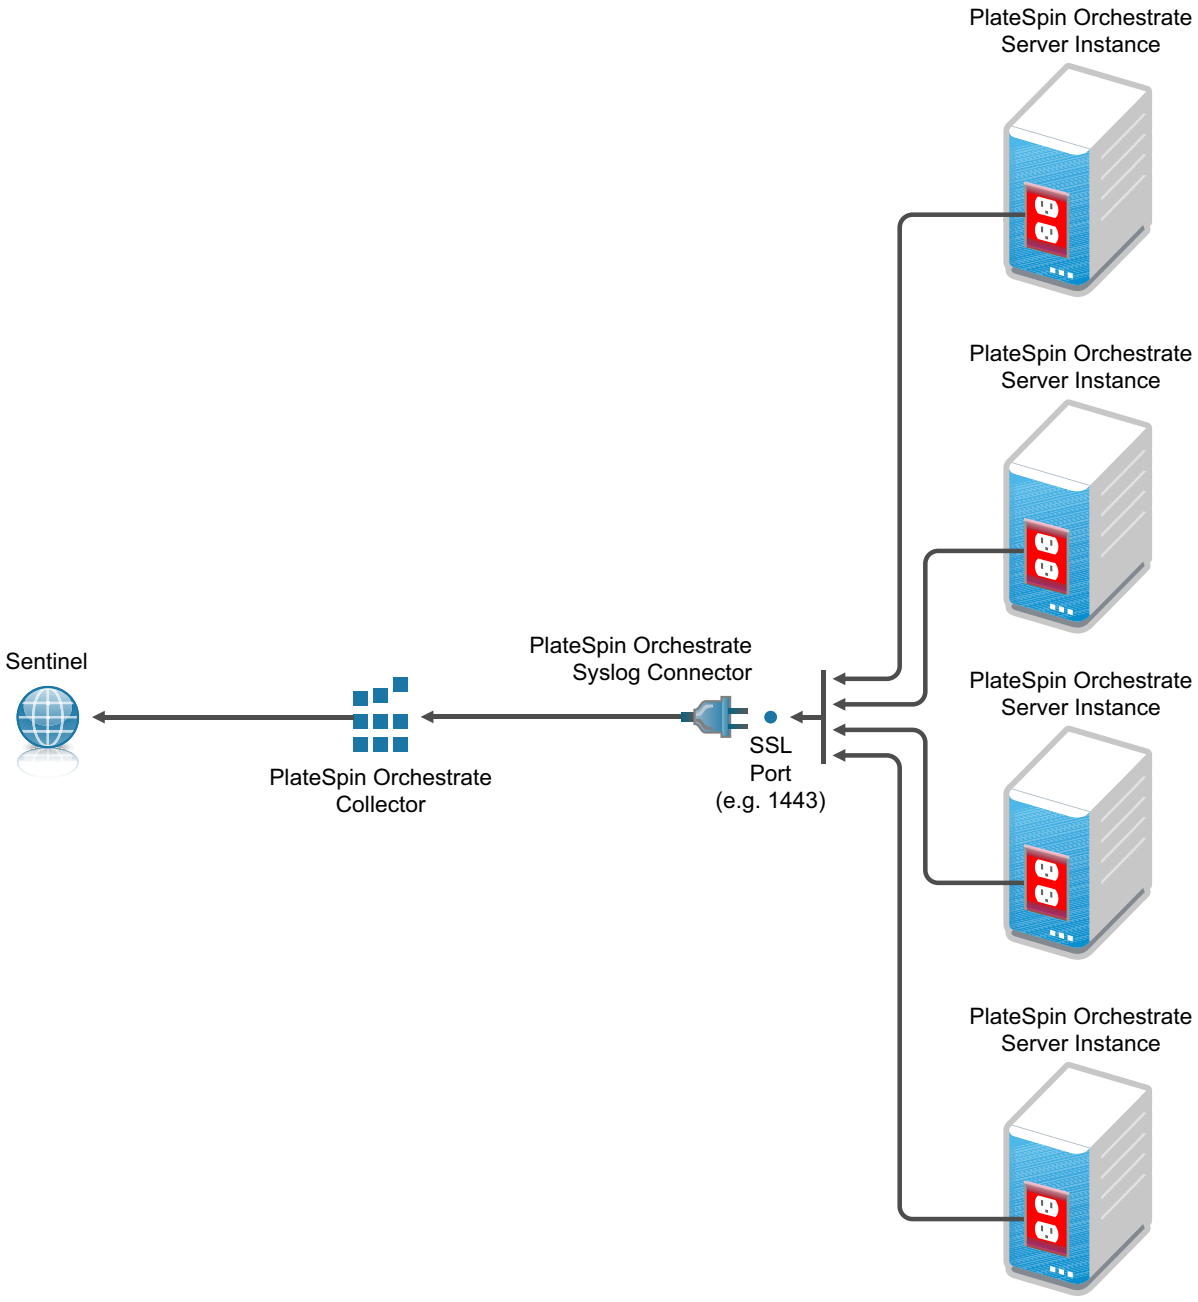

NOTE: Multiple Orchestration Server instances can send syslog messages to a single Syslog Connector.

17.2 System Requirements

Integrating Sentinel and the Orchestration Server requires the following:

- Cloud Manager Orchestration Server 3.5 installed and running on a supported SUSE Linux Enterprise Server.

- Sentinel 6.1.1.1 or later installed and running. For information about Sentinel, see the NetIQ Sentinel How to Buy (https://www.netiq.com/products/sentinel/how-to-buy/) page.

  NOTE: Sentinel must be separately purchased from NetIQ.

  Sentinel can be installed on the same server with the Orchestration Server, assuming that the server has sufficient RAM, processing power, and the disk space to accommodate running the two products side-by-side; otherwise they should be run on separate servers that communicate through TCP/IP.

  For more information about the server requirements for the Orchestration Server, see Section 2.2, “Cloud Manager Orchestration Server Requirements,” on page 18.

  For more information about the server requirements for Sentinel, see the Technical Information for NetIQ Sentinel (https://www.netiq.com/Support/sentinel/techinfo.asp) page.

- The appropriate Sentinel Syslog Connector, available for download from the NetIQ Sentinels Plug-ins > Connectors (https://www.netiq.com/support/sentinel/plugins/) page.
- The appropriate Orchestration Server/Sentinel Collector Plug-in, available for download on the NetIQ Sentinels Plug-ins > Collectors (https://www.netiq.com/support/sentinel/plugins/) page.

  NOTE: Collector-specific information is available in `.pdf` format with the download above. Users should consult with NetIQ Support to determine the appropriate connector and collector for use with the Orchestration Server.

17.3 Importing and Deploying the Orchestration Server Sentinel Collector Plug-in

1. From the Sentinel Control Center, click **Event Source Management > Live View** to display the Event Source Management (Live View) window.
2. From the Event Source Management (Live View) window, click **Tools > Import plug-in** to display the Import Plug-in Wizard.
3. From the Import Plug-in Wizard, select **Import Collector Script or Connector plugin package file (.zip)**, then click **Next** to open a file browser.

   IMPORTANT: When you try to import the syslog connector, you might see a message like this:

   ```
   No universal SYSLOG collector found. Resolve this problem by importing either the NetIQ Generic Collector or another Syslog collector that is designated as a universal Syslog collector.
   ```

   You can safely ignore this message. If you need more information, see the *Syslog Connector Installation Guide*.

4. From the file browser, locate and select the Orchestration Server Collector `.zip` file, `NetIQ_Cloud-Manager-Orchestration-Service_6.1r.clz.zip` (or similar), then click **OK** to display the Plugin Details page.
5. On the Plugin Details page, click **Finish** to display the Graph page of the Event Source Management Live View.
6. On the Graph page, right-click the Collector Manager icon, then click **Add Event Source Server** to display the Select Connector Plugin page of the Add Event Source Server Wizard.

7 From the Select Connector Plugin page, ensure that the Syslog connector is selected in the Installed Connectors table, then click **Next** to display Syslog Event Source Server page of the wizard.

8 On the Syslog Event Source Server page of the tool, Select **SSL**, enter `1443` as the default port number, then click **Next** to display the **Message Handling** page of the wizard.

Port 1443 is used in this example to eliminate conflict with other NetIQ products.

It is not required that you use `1443` in this field, but any port number that you configure here must match the port number you assign for Sentinel integration in the Orchestration Console. For more information, see "Sentinel Server Configuration Panel" in the *Cloud Manager Administrator Reference*.

9 Click **Next** on the Message Handling page and succeeding pages of the wizard until you reach the General page.

10 On the **General** page of the wizard, ensure that the **Run** check box is selected, then click **Finish** to display the revised Graph page of the Event Source Management Live View.

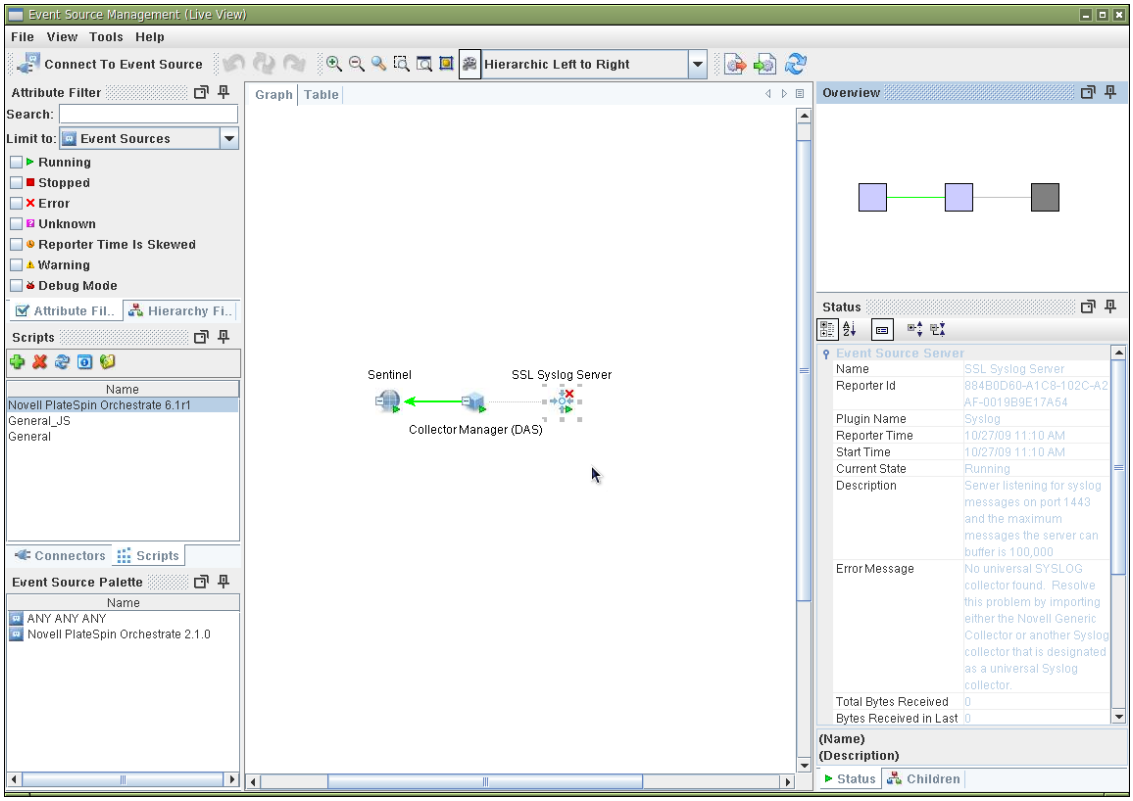

In this graph page, the Event Source Server icon displays a superimposed red cross. The view also displays a warning (`No universal SYSLOG connector found...`) is visible in the right hand pane. You can safely disregard these warnings.

17.4 Connecting the Orchestration Server to the Sentinel Collector Plug-In

After you deploy the Orchestration Server/Sentinel Collector Plug-in, use the following steps to configure the connection between the plug-in and the Orchestration Server.

1. From a running Orchestration Console, select the Grid Server object that you want to connect to the Sentinel Collector Plug-in, then log in to that grid.
2. In Admin View of the object, select **Info/Configuration**, then scroll to the **Sentinel Server Configuration** panel in this view to configure the collector.
3. In the **Server Hostname** and **Server Port Number** fields, ensure that the server host name points to the host running the Sentinel Event Source Server and that the port value is listed as `1443`, then click **Connect**.
4. Verify that the **Is Connected** check box is selected.

For more information about these fields in the Orchestration Console Admin View, see "Sentinel Server Configuration Panel" in the *Cloud Manager Administrator Reference*.

17.5 Verifying the Sentinel Configuration After Connecting to the Orchestration Server

The first time that the Orchestration Server connects to Sentinel through the SSL port, Sentinel automatically creates and configures its Connector, Collector, and Event Source Server. You can verify this on the Graph page of the Event Source Management Live View.

Figure 17-2 *Event Source Server Diagram Showing the Orchestration Server and Sentinel Server Connections*

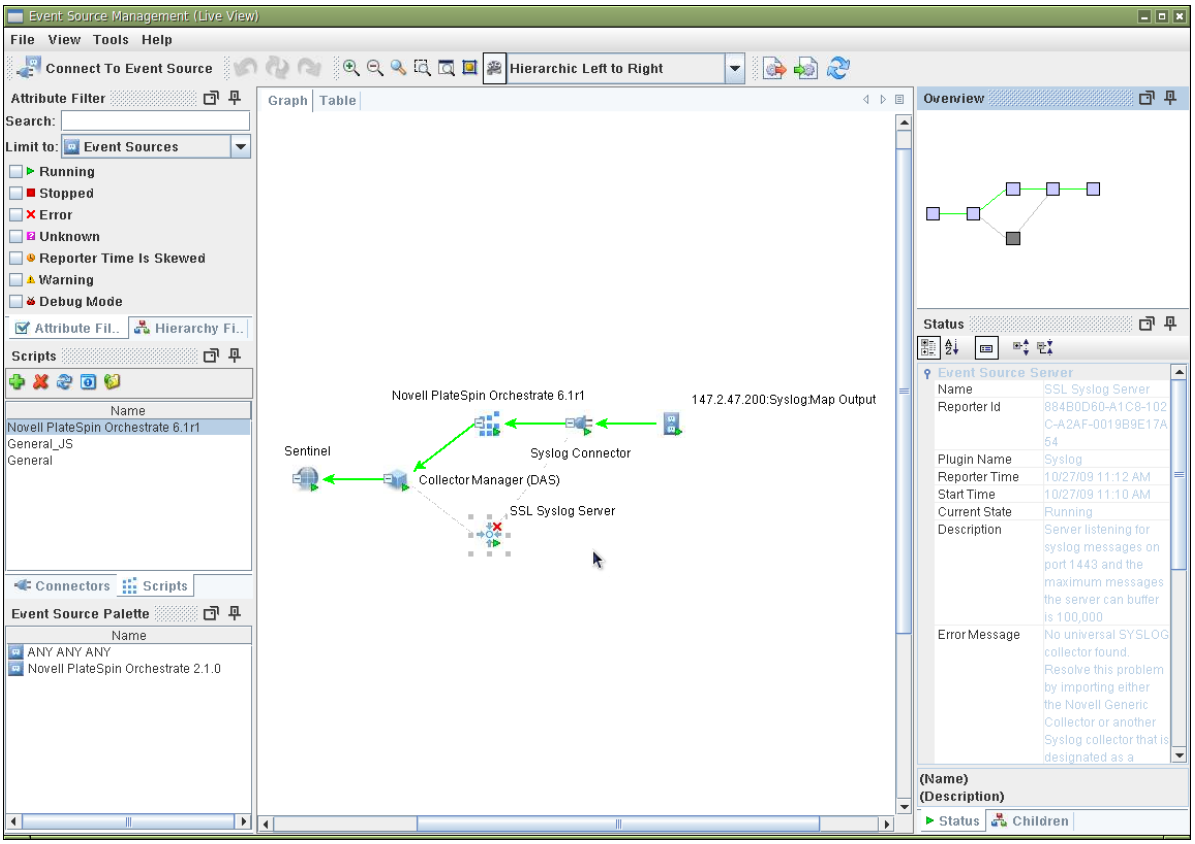

17.6 Event Classification and Taxonomy Keys

The value of the `eventclass` key in Sentinel's **ExtendedInformation** event field specifies a class of events where this particular event belongs. This mechanism allows classification of multiple events into a single event class, even when their message bodies and/or other `ExtendedInformation` key/value pairs differ. For example, any event detailing the logout of a user from an Orchestration Server administrator account is classified with the event class `admin_logout`. This event class also serves as the taxonomy key for use by Sentinel.

Many log events are currently unclassified, so their event classes are also unclassified. We anticipate the number of unclassified log events to decrease incrementally over subsequent Cloud Manager Orchestration Server releases.

All non-failure events are logged on success rather than on action initiation. For example, an event with the event class `job_deployed` would only be seen after the job had been successfully deployed, not when the deployment attempt was initiated.

Table 17-1 *eventclass Taxonomy Keys*

Value of the eventclass Taxonomy Key	Events in This Classification	Other Keys Typically Used with This Class of Events
`unclassified`	Any event not yet assigned an event class.	

Value of the eventclass Taxonomy Key	Events in This Classification	Other Keys Typically Used with This Class of Events
sentinel	Events relating to integration of Cloud Manager Orchestration Server with Sentinel.	
authorization_failure	Any kind of authorization failure.	action
authentication_failure	Any kind of authentication failure.	action
admin_login	Login to a Orchestration Server administrator account (for example, through zosadmin or through the Orchestration Console).	user
admin_logout	Logout from a Orchestration Server administrator account (for example, through zosadmin or through the Orchestration Console).	user
user_login	Logout from a grid user account.	user
user_logout	Login of a resource to the grid.	user
user_password_change	Password change for grid user account.	user
resource_login	Logout of a resource from the grid.	resource
resource_logout	Logout of a resource from the grid.	resource
repository_created	New repository created.	repository
repository_deleted	Repository deleted.	repository
resource_created	New resource created.	resource, type
resource_deleted	Resource deleted.	resource
user_created	New user created.	user
user_deleted	User deleted.	user
session_not_found	User or agent session not found.	session
vbridge_created	New vBridge created.	vbridge, vmhost
vbridge_deleted	vBridge deleted.	vbridge
vdisk_created	New vDisk created.	vdisk, vm
vdisk_deleted	vDisk deleted.	vbridge
vnic_created	New vNIC created.	vnic, vm
vnic_deleted	vNIC deleted.	vnic
group_created	New group created.	group, type
group_deleted	Group deleted.	group, type
group_member_added	New member added to the Grid object group.	group, member, type
group_member_removed	Member removed from the Grid object group.	group, member, type
job_deployed	New job deployed to the grid.	job
job_undeployed	Job undeployed from the grid.	job

Value of the eventclass Taxonomy Key	Events in This Classification	Other Keys Typically Used with This Class of Events
schedule_deployed	New schedule deployed to the grid.	schedule
schedule_undeployed	Schedule undeployed from the grid.	schedule
trigger_deployed	New trigger deployed to the grid.	trigger
trigger_undeployed	Trigger undeployed from the grid.	trigger
policy_association	Association of an Orchestration Server policy with a Grid object.	policy, target
policy_disassociation	Disassociation of an Orchestration Server policy with a Grid object.	policy, target
policy_created	Policy added to the grid.	policy
policy_removed	Policy removed from the grid.	policy
job_started	Job (instance) started.	jobinstance
job_finished	Job (instance) finished.	jobinstance, outcome (completed, canceled, or failed), reason (when canceled)
vm_applyconfig	VM apply config action initiated.	vm, user
vm_build	VM build action initiated.	vm, user
vm_check_status	VM check status action initiated.	vm, user
vm_checkpoint	VM checkpoint action initiated.	vm, user, checkpoint
vm_clone	VM clone action initiated.	vm, user
vm_create_template	VM create_template action initiated.	vm, user
vm_delete	VM delete action initiated.	vm, user
vm_destroy	VM destroy action initiated.	vm, user
vm_install_agent	VM install agent action initiated.	vm, user
vm_make_standalone	VM make standalone action initiated.	vm, user
vm_migrate	VM migrate action initiated.	vm, user
vm_move	VM move action initiated.	vm, user
vm_pause	VM pause action initiated.	vm, user
vm_personalize	VM personalize action initiated.	vm, user
vm_provision	VM provision action initiated.	vm, user
vm_restart	VM restart action initiated.	vm, user
vm_restore	VM restore action initiated.	vm, user, checkpoint
vm_resume	VM resume action initiated.	vm, user
vm_saveconfig	VM save config action initiated	vm, user
vm_shutdown	VM shutdown action initiated.	vm, user

Value of the eventclass Taxonomy Key	Events in This Classification	Other Keys Typically Used with This Class of Events
vm_suspend	VM suspend action initiated.	vm, user
resource_healthy	The Resource became healthy.	resource
resource_unhealthy	The Resource became unhealthy.	resource
repository_healthy	The Repository became healthy.	repository
repository_unhealthy	The Repository became unhealthy.	repository
vbridge_healthy	The vBridge became healthy.	vbridge
vbridge_unhealthy	The vBridge became unhealthy.	vbridge
vnic_healthy	The vNIC became healthy.	vnic
vnic_unhealthy	The vNIC became unhealthy.	vnic
vdisk_healthy	The vDisk became healthy.	vdisk
vdisk_unhealthy	The vDisk became unhealthy.	vdisk
vmhost_healthy	The VM host became healthy.	vmhost
vmhost_unhealthy	The VM host became unhealthy.	vmhost

17.7 Plain Text Visibility of Sensitive Information

The following table outlines where sensitive information might be visible as plain text:

Table 17-2 *Locations Where Sensitive Information Might Be Stored As Plain Text*

Information	Storage Location	Visibility Issue
Audit Database configuration	ZOS properties store	Contains plain text information including user/password for allowing the Orchestration Server to log into the Audit Database for logging. You should use a non-privileged database account for logging.

18 Configuring Secure Authentication Sources to Communicate with Cloud Manager

This section discusses configuring NetIQ Access Manager (NAM) as identity service tools that you can leverage to let your Cloud Manager users securely log in to Cloud Manager.

- Section 18.1, “Configuring NetIQ Access Manager to Work with Cloud Manager,” on page 171

18.1 Configuring NetIQ Access Manager to Work with Cloud Manager

NetIQ Access Manager (NAM) provides secure, single sign-on access to trusted Cloud Manager users from any location, in spite of the internal technical and organizational boundaries in your enterprise. NetIQ Access Manager supports multi-factor authentication, role-based access control, data encryption, and SSL VPN services.

The content in this section is not intended as a comprehensive guide to NAM. You should have already installed NetIQ Access Manager and a NetIQ Access Manager Access Gateway. You should also have installed Cloud Manager, and the Cloud Manager Application Server should be running.

You need to be familiar with NetIQ Access Manager capabilities so that you understand the context of the content in this section. For more information about NetIQ Access Manager, see the Access Manager documentation website (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/).

- Section 18.1.1, “Managing a Reverse Proxy for Authentication to Cloud Manager,” on page 171

18.1.1 Managing a Reverse Proxy for Authentication to Cloud Manager

A reverse proxy acts as the front end to the Cloud Manager Web Server on your Internet. The proxy off-loads frequent requests, thereby freeing up bandwidth. It also increases security because the IP addresses of your Web servers are hidden from the Internet.

You can use an existing reverse proxy and add a new proxy service for Cloud Manager or you can create a new reverse proxy with a service for Cloud Manager. You can configure the authentication settings of the reverse proxy according to the needs of your enterprise.

For information about creating a new reverse proxy, see *“Managing Reverse Proxies and Authentication”* (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/accessgatewayhelp/data/reverselist.html) in the *NetIQ Access Manager Appliance 4.0 Access Gateway Guide*.

When the reverse proxy is set up as you want it, you need to perform the other configuration procedures necessary for NetIQ Access Manager authentication:

- “Creating and Configuring the Proxy Service for the Cloud Manager Reverse Proxy” on page 172
- “Adding and Protecting All Cloud Manager Resources” on page 172
- “Creating an Identity Injection Policy for the New Cloud Manager Protected Resource” on page 174
- “Adding and Configuring an HTML Rewriter Profile for the Proxy Service” on page 176

Creating and Configuring the Proxy Service for the Cloud Manager Reverse Proxy

You must create a unique proxy service for Cloud Manager. Configure the proxy service settings according to the needs of your enterprise.

The first proxy service of a reverse proxy is considered the master (or parent) proxy. Subsequent proxy services can use domain-based, path-based, or virtual multi-homing, relative to the published DNS name of the master proxy service. If you are creating a second proxy service to be used for Cloud Manager on the reverse proxy, see “*Using Multi-Homing to Access Multiple Resource”s (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/accessgatewayhelp/data/b34l8ue.html)* in the *NetIQ Access Manager Appliance 4.0 Access Gateway Guide*.

Remember that for the **Web Server IP Address** setting of the proxy service, you need to specify the IP Address for the Cloud Manager Web server, and for the **Web Server Host Name** setting of the proxy service, you need to specify the DNS name of the Cloud Manager Web server.

When you have configured the proxy service according to your needs, you can continue with “Adding and Protecting All Cloud Manager Resources” on page 172.

Adding and Protecting All Cloud Manager Resources

A protected resource configuration specifies the directory (or directories) on the Cloud Manager Web server that you want to protect. The protected resource configuration specifies the authorization procedures and the policies that should be used to enforce protection.

You need to group all of the Cloud Manager resources that use the proxy service.

To create a resource that groups all of the Cloud Manager services:

1. Log in to the Access Manager Administration Console. For information about accessing the console, see *“Logging In to the Administration Console”* (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/target_installation/data/b139b2ny.html#b139d8i3) in the *NetIQ Access Manager Appliance 4.0 Installation Guide*.
2. In the console, select **Devices** > **Access Gateways** to open the Access Gateways page.
3. On the Access Gateways page, select **Edit** for the Gateway Server you want to edit. This displays the Access Gateway Server Configuration page.
4. On the Access Gateway Server Configuration page, select the name of the reverse proxy. This opens the Reverse Proxy configuration page.
5. On the Reverse Proxy page, select the proxy service you want to configure. This opens the Reverse Proxy Service page.
6. On the Reverse Proxy Service page, select the **Protected Resources** tab to open the Protected Resources page.



7 Configure the protected resource.:

7a On the Protected Resources page, select **New**, then specify a display name for the new resource you want to protect. For example, to create a resource that you want to use to represent all Cloud Manager resources, you could name the resource “everything.”

When you create the display name, the Overview page for the new resource is displayed.

7b Fill in the fields to configure the resource:

- **Description:** Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.
- **Authentication Procedure:** Select **Name/Password -Form** from the drop-down list. This specifies a form-based authentication over HTTP or HTTPS, using the Access Manager login form.
- **URL Path:** Select the default path, which is `/*`. This specifies everything on the Cloud Manager Web Server.

7c Click the **Protected Resources** breadcrumb at the top of the Overview page to return to the Protected Resources page.

7d On the Protected Resources page, ensure that the new protected resource is selected as **Enabled**.

8 Continue with “Creating an Identity Injection Policy for the New Cloud Manager Protected Resource” on page 174.

Creating an Identity Injection Policy for the New Cloud Manager Protected Resource

When the Cloud Manager protected resource is created, you need to associate it with an Access Manager identity injection policy to protect it. This policy specifies the information that must be injected into the HTTP header. Because Cloud Manager is configured to detect certain fields in the header, it can deny user authentication or redirect that user to an alternate web page if it does not find the required information in the header.

1. Log in to the Access Manager Administration Console. For information about accessing the console, see *"Logging In to the Administration Console"* (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/target_installation/data/b139b2ny.html#b139d8i3) in the *NetIQ Access Manager Appliance 4.0 Installation Guide*.
2. In the Access Manager Administration Console, select **Devices** > **Access Gateways** to open the Access Gateways page.
3. On the Access Gateways page, select **Edit** for the Gateway Server you want to edit. This displays the Access Gateway Server Configuration page.
4. On the Access Gateway Server Configuration page, select the name of the reverse proxy. This opens the Reverse Proxy configuration page.
5. On the Reverse Proxy page, select the proxy service you want to configure. This opens the Reverse Proxy Service page.
6. On the Reverse Proxy Service page, select the **Protected Resources** tab to open the Protected Resources page.
7. On the Protected Resources page, select the display name of the Cloud Manager protected resource to open the properties views, then select **Identity Injection** to open the Identity Injection Policy List.
8. Select **Manage Policies** to open the Policies page.
9. Fill in the fields.
   - **Description:** (Optional) Describe the purpose of this policy. Because Identity Injection policies are customized to match the content of a specific Web server, include the name of the Cloud Manager Web server as part of the description.
   - **Priority:** Specify the order in which a rule is applied in the policy, when the policy has multiple rules. The highest priority is 1 and the lowest priority is 10.
10. In the actions panel of the page, select **New** > **Inject into Custom Header**.

    This inserts custom names with values into a custom header.
11. Configure five custom policy headers for Cloud Manager. You must configure the attributes of the custom headers as specified below. The headers must be created or moved into the order listed. You can use the **Copy Action** icon to copy each header, then you can modify the configurations as needed.

    11a Create the `X-TrustedUser` header, using the following information to populate the fields.:
    - **Custom Header Name:** Specify `X-TrustedUser`.
    - **Value:** Select **LDAP Attribute**. Selecting this option enables the LDAP attribute list box and the Refresh Data Rate list box. For this header, select **cn** as the LDAP attribute, then select **Session** as the refresh rate.
    - **Multi-Value Separator:** Select the semicolon (;) separator from the list box.
    - **DN Format:** Select the **LDAP** option from the list box.

11b Create the `X-TrustedRoles` header, using the following information to populate the fields:

- **Custom Header Name:** Specify `X-TrustedRoles`.
- **Value:** Select **LDAP Attribute**. Selecting this option enables the LDAP attribute list box and the Refresh Data Rate list box. For this header, select **groupMembership** as the LDAP attribute, then select **Session** as the refresh rate.

  NOTE: The **groupMembership** attribute applies if you are using eDirectory. If you are using Active Directory, the attribute is **memberOf**.
- **Multi-Value Separator:** Select the semicolon (;) separator from the list box.
- **DN Format:** Select the **LDAP** option from the list box.

11c Create the `X-TrustedUserFQDN` header, using the following information to populate the fields:

- **Custom Header Name:** Specify `X-TrustedUserFQDN`.
- **Value:** Select **Credential Profile**. Selecting this option enables the Credential Profile list box. For this header, select **LDAP Credentials: LDAP User DN** as the credential profile.
- **Multi-Value Separator:** Select the semicolon (;) separator from the list box.
- **DN Format:** Select the **LDAP** option from the list box.

11d Create the `X-TrustedUserDisplayName` header using the following information to populate the fields.

- **Custom Header Name:** Specify `X-TrustedUserDisplayName`.
- **Value:** Select **LDAP Attribute**. Making this selection enables the LDAP attribute list box and the Refresh Data Rate list box. For this header, select **displayName** as the LDAP attribute, then select **Session** as the refresh rate.
- **Multi-Value Separator:** Select the semicolon (;) separator from the list box.
- **DN Format:** Select the **LDAP** option from the list box.

11e Create the `X-TrustedUserEmail` header using the following information to populate the fields.

- **Custom Header Name:** Specify `X-TrustedUserEmail`.
- **Value:** Select **LDAP Attribute**. Making this selection enables the LDAP attribute list box and the Refresh Data Rate list box. For this header, select **mail** as the LDAP attribute, then select **Session** as the refresh rate.
- **Multi-Value Separator:** Select the semicolon (;) separator from the list box.
- **DN Format:** Select the **LDAP** option from the list box.

12 Click **OK** to save the new policy and display it on the Policies page.

13 On the Policies page, click **Enable** to enable this new policy for the protected resource.

14 Continue with "Adding and Configuring an HTML Rewriter Profile for the Proxy Service" on page 176.

NOTE: Ensure that you always update your configuration when you make changes in NetIQ Access Manager.

For more information, see *Creating Identity Injection Policies* (https://www.netiq.com/documentation/netiqaccessmanager4/policyhelp/data/b5547ku.html) in the *NetIQ Access Manager 4.0 Policy Guide*.

Adding and Configuring an HTML Rewriter Profile for the Proxy Service

The changes you make to the NetIQ Access Manager Access Gateway configurations for Cloud Manager require HTML rewriting because the Cloud Manager Web server is not aware that the Access Gateway machine is obfuscating its DNS names. URLs contained in its pages must be checked to ensure that these references contain the DNS names that the client browser understands. On the other end, the client browsers are not aware that the Access Gateway is obfuscating the DNS names of the resources they are accessing. The URL requests coming from the client browsers that use published DNS names must be rewritten to the DNS names that the Cloud Manager Web server expects.

The information in *Configuring HTML Rewriting (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/accessgatewayhelp/data/b3nqotc.html)* in the *NetIQ Access Manager Appliance 4.0 Access Gateway Guide* explains this process more fully.

You need to create and configure a new HTML Rewriter Profile for use with Cloud Manager.

1. Log in to the Access Manager Administration Console. For information about accessing the console, see *"Logging In to the Administration Console"* (https://www.netiq.com/documentation/netiqaccessmanager4_appliance/target_installation/data/b139b2ny.html#b139d8i3) in the *NetIQ Access Manager Appliance 4.0 Installation Guide.*
2. In the Access Manager Administration Console, select **Devices** > **Access Gateways** to open the Access Gateways page.
3. On the Access Gateways page, select **Edit** for the Gateway Server you want to edit. This displays the Access Gateway Server Configuration page.
4. On the Access Gateway Server Configuration page, select the name of the reverse proxy. This opens the Reverse Proxy configuration page.
5. On the Reverse Proxy page, select the proxy service you want to configure. This opens the Reverse Proxy Service page.
6. On the Reverse Proxy Service page, select the **HTML Rewriting** tab to open the HTML rewriting page.

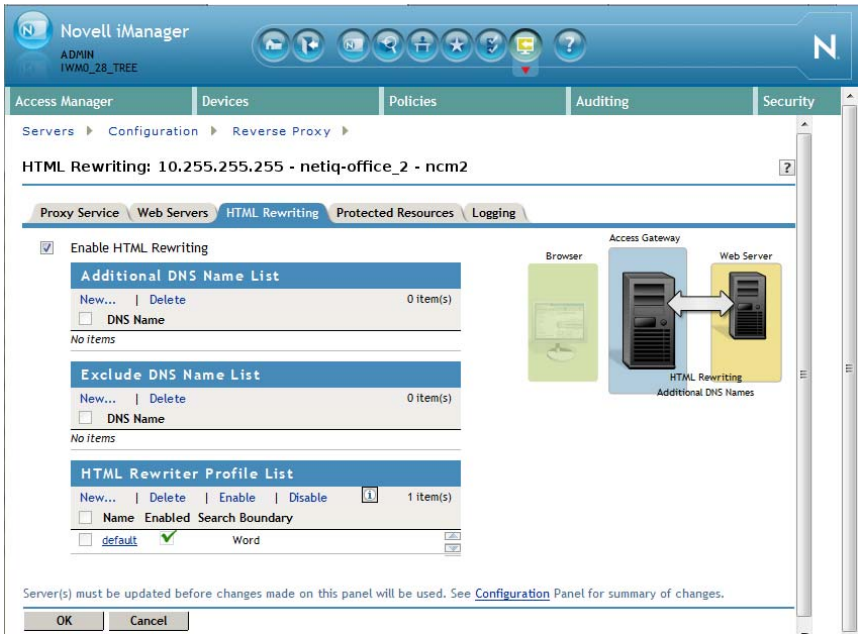

The HTML Rewriting page specifies which DNS names are to be rewritten. The HTML Rewriter Profile specifies which pages to search for DNS names that need to be rewritten.

7 Select **Enable HTML Rewriting**.

This option is enabled by default. When it is disabled, no rewriting occurs. When it is enabled, this option activates the internal HTML rewriter. When data is sent to the browsers, this rewriter replaces the name of the Cloud Manager Web server with the published DNS name. It replaces the published DNS name with the Web Server Host Name when sending data to the Cloud Manager Web server. It also ensures that the proper scheme (HTTP or HTTPS) is included in the URL. This is needed because you can configure the Access Gateway to use HTTPS between itself and client browsers and to use HTTP between itself and the Web servers.

8 Specify a name for the new profile, use the default search boundary, then click **OK** to open the HTML Rewriter configuration page.

9 In the **Content-Type Header** section of the page, click **New** to open a New dialog.

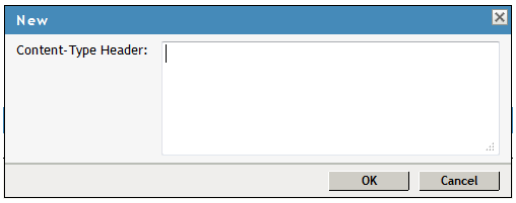

10 In the dialog, specify the new content-type header, which is `application/xml`, select the **Rewrite Inbound Headers** check box, then click **OK** to ensure that the new Content-Type Header is enabled for the protected resource.

V Upgrading Cloud Manager

This section includes information about upgrading the Cloud Manager components:

- Chapter 19, “Upgrade Checklist,” on page 181
- Chapter 20, “Upgrading the Cloud Manager Application Server Components,” on page 183
- Chapter 21, “Orchestration Components Upgrade Overview,” on page 187
- Chapter 22, “Upgrading Orchestration Server Components,” on page 189
- Appendix A, “Restoring Cloud Manager Application Server in the Event of a System Failure,” on page 205
- Appendix B, “Compatibility Checking Behavior for Orchestration Components,” on page 207
- Appendix C, “Recovering from a Failed Orchestration Server Upgrade,” on page 211

19 Upgrade Checklist

Use the upgrade checklist in this section to plan your Cloud Manager upgrade.

Status	Preparation and Upgrade Tasks	For information
❑	1. Back up the CMAS cloudmanager database instance.	Section 20.1, “Backing Up the PostgreSQL Database,” on page 183
❑	2. (Optional) Back up the CMAS components, including custom files you might have created for it.	Section 20.2, “Backing Up the Application Server Components and Custom Files,” on page 184
❑	3. Install the upgrade package for Application Server.	Section 20.3, “Install the Cloud Manager Package,” on page 185
❑	4. Configure the updated Application Server.	Section 20.4, “Running the Cloud Manager Configuration Script,” on page 185
❑	5. Ensure that you understand the planning requirements for the CMOS upgrade.	Chapter 21, “Orchestration Components Upgrade Overview,” on page 187
❑	6. Back up the CMOS components.	Section 22.1, “Backing Up Orchestration Components Prior to Upgrading,” on page 189
❑	7. Check the current version of CMOS components running in your Cloud Manager system.	Section 22.2, “Checking the Current Version of Cloud Manager Orchestration Components,” on page 190
❑	8. Ensure that all running jobs are complete and take a snapshot of the current configuration of the Orchestration Server.	Section 22.3, “Snapshotting the Existing Orchestration Server Installation,” on page 190
❑	9. Install the upgrade packages for Orchestration components: ◆ Orchestration Server ◆ Orchestration Console ◆ Orchestration Agent ◆ Monitoring Server ◆ Monitoring Agent	Section 22.4, “Upgrading Orchestration Components,” on page 191
❑	10. Configure the updated Orchestration components.	Section 22.5, “Configuring the Upgraded Orchestration Components,” on page 195
❑	11. Configure the remote audit database.	Section 22.6, “Configuring the Remote Audit Database after Orchestration Components Are Upgraded,” on page 199
❑	12. Run Discovery on VM Hosts and Images.	Section 22.7, “Running Discovery on VM Hosts and Images,” on page 200
❑	13. Update resource pool data on the Cloud Manager Application Server.	Section 20.5, “Using the Configurator Tool to Update Resource Pool Data on the Cloud Manager Application Server,” on page 186

20 Upgrading the Cloud Manager Application Server Components

This section includes information you need for upgrading the Cloud Manager Application Server 2.4 to Cloud Manager Application Server 2.5. It is important that you upgrade the Application Server in the sequence that follows:

- Section 20.1, “Backing Up the PostgreSQL Database,” on page 183
- Section 20.2, “Backing Up the Application Server Components and Custom Files,” on page 184
- Section 20.3, “Install the Cloud Manager Package,” on page 185
- Section 20.4, “Running the Cloud Manager Configuration Script,” on page 185
- Section 20.5, “Using the Configurator Tool to Update Resource Pool Data on the Cloud Manager Application Server,” on page 186

For information about upgrading the Cloud Manager Orchestration Server, see Chapter 22, “Upgrading Orchestration Server Components,” on page 189.

20.1 Backing Up the PostgreSQL Database

As a precaution, you should always back up the PostgreSQL database that you use with Cloud Manager before you upgrade to a newer version of Cloud Manager.

For more information about backing up a PostgreSQL database, see the PostgreSQL documentation (http://www.postgresql.org/docs/9.1/static/backup.html). The following specific instructions are helpful in this process:

- If you chose to automatically configure the PostgreSQL database during the initial Cloud Manager installation, the PostgreSQL password is located in the `/etc/opt/netiq/cloudmanager/etc/pgusr.in` file.
- If the PostgreSQL DBMS you are using has several databases, you can list the database names using the following command:

  ```
  sudo -u postgres psql -l
  ```

- The Cloud Manager database name should contain the phrase `cloudmanager`. This is the database you want to back up.
- When you identify the Cloud Manager database you want to back up, use a command similar to the following to back it up:

  ```
  sudo -u postgres pg_dump -s cloudmanager > /var/opt/netiq/cloudmanager/
  cm_database_backup.dump.out
  ```

20.2 Backing Up the Application Server Components and Custom Files

Before you upgrade the Cloud Manager Application Server, you can back up the Application Server components in the current Cloud Manager system and your custom files. Cloud Manager 2.5 includes a shell script tool that lets you back up the current Cloud Manager 2.4 system prior to the upgrade so that if the upgrade to 2.5 fails, you can revert to the 2.4 configuration settings and files. The `backup` tool backs up the CMAS component files (with the exception of the PostgreSQL database) and any custom files that you might have created for it. This is an optional process; normally, the backup of the Cloud Manager configuration that is automatically triggered during an upgrade is sufficient for Cloud Manager 2.4 restoration.

NOTE: A reversion to Cloud Manager 2.4 settings and files pre-supposes that you would revert to an earlier version of the PostgreSQL database.

To perform the pre-upgrade backup of CMAS components:

1. Download the appropriate Cloud Manager ISO, then find and extract the `backup` tool.

   The upgrade tool is located in the `/tools` directory at the root of the ISO. Ensure that you extract it to the computer where your current Cloud Manager 2.4 system resides.

2. Run the `backup` tool, using the appropriate parameters.

   Use the table below to help you understand the options you must use with the tool and those that are optional.

Backup Tool Parameter	Function
Required Parameters (You must use either of the following)	
`-b`	Perform a backup of the files and directories.
`-r`	Perform a restore of the files and directories.
`-f <backup_file>`	Create a backup of all necessary Cloud Manager files and folders in the specified compressed file. If this option is specified along with the restore flag (-r), it extracts the files from the <*backup_file*> into their proper location
Optional Parameters	
`-c`	Back up Cloud Manager configuration files only.
`-s`	Perform a “silent” backup, then print the final backup location to the terminal.

   Syntax: Structure the backup command like this:

   `./backup < required_parameters > < optional_parameters >`

   Example 1: Run the tool to *create* a backup of a Cloud Manager 2.4 system like this:

   ```
   ./backup -b -f /var/opt/netiq/cloudmanager/my_backup_name.tar.gz
   ```

   Example 2: Run the tool to *restore* a backup of a Cloud Manager 2.4 system like this:

   ```
   ./backup -r -f /var/opt/netiq/cloudmanager/my_backup_name.tar.gz
   ```

20.3 Install the Cloud Manager Package

1. Log in as the `root` user on your Cloud Manager 2.4 Application Server.
2. Install the Cloud Manager 2.5 RPM files using either of the following methods:
   - Use YaST2 as described in Section 6.1, “Installing Cloud Manager Pattern on Your SLES Server,” on page 55.
   - Use the following command at the Application Server command line:

     ```
     rpm -Uvh netiq-cloudmanager-2.5.0-<build_number>.noarch.rpm
     ```
3. Continue with Section 20.4, “Running the Cloud Manager Configuration Script,” on page 185.

   Do not start Cloud Manager Services until the configuration process is complete.

20.4 Running the Cloud Manager Configuration Script

After you have successfully installed the new Cloud Manager 2.5 package, you must configure that package by running the Cloud Manager configuration script. The upgrade installer saves all current configuration data. You can use this configuration data to facilitate a system recovery.

To configure the updated Application Server packages:

1. At the command line, run the following command:

   ```
   /opt/netiq/cloudmanager/configurator/config
   ```

   The configurator script recognizes that a new version of Cloud Manager is installed and prompts with the following text:

   ```
   Welcome to the NetIQ Cloud Manager configuration utility.

       One or more products on this system require an upgrade.

       Select whether to perform the required upgrade to an existing configuration
       (upgrade), or to run the new install configuration (install) for a product.

   u) upgrade
   i) install
   - - - - - -

   Selection [upgrade]:
   ```
2. Specify `u` to indicate that this is an upgrade for Cloud Manager, then press Enter.
3. Specify the name of the PostgreSQL administrator, then press Enter.

   This must be the same PostgreSQL user name that you specified in the previous Cloud Manager installation.
4. Specify the password for the PostgreSQL administrator, then press Enter.

   This must be the same PostgreSQL password that you specified in the previous Cloud Manager installation.
5. (Conditional, if you have backed up the PostgreSQL database, as explained in “Backing Up the PostgreSQL Database” on page 183). Specify `yes` to indicate that you have backed up the PostgreSQL database, then press Enter.
6. After the configuration summary is displayed, specify `yes` to start the configuration process.
7. Start Cloud Manager Services:

   ```
   /etc/init.d/netiq/cloudmanager start
   ```

You can start Cloud Manager Services when the configuration process is complete.

NOTE: If you start Cloud Manager services after you have upgraded the package but before you have configured the upgrade, Cloud Manager displays an error.

The settings you specify in the configuration file are recorded in `/etc/opt/netiq/cloudmanager/netiq_cloudmanager_config.conf`. You can find a log file of the actual configuration process at `/var/opt/netiq/cloudmanager/logs/netiq_cloudmanager_config.log`.

After the Cloud Manager Application Server 2.4 is upgraded to Cloud Manager Application Server 2.5, you can roll back to Cloud Manager Application Server 2.4 if failures occur on the upgraded system. See Appendix A, “Restoring Cloud Manager Application Server in the Event of a System Failure,” on page 205.

20.5 Using the Configurator Tool to Update Resource Pool Data on the Cloud Manager Application Server

After all Orchestration grids (also known as Cloud Manager Zones) are upgraded, you can use the Cloud Manager Application Server configurator utility to perform a VMware resource pool ID data upgrade. If an Orchestration grid (zone) has not been updated or if the grid is inaccessible to the Cloud Manager Application Server, the IDs stored in the Application Server for the resource pools belonging to that Cloud Manager Zone are not upgraded. You can upgrade the data later, if you choose to, when the zone has been updated with the Orchestration upgrade.

We recommend that you upgrade your Orchestration Server grids (that is, Cloud Manager “zones”) as soon as possible in order to avoid unexpected behavior.

To run the Application Server upgrade for resource pool data:

1. Ensure that the Cloud Manager service is running.
2. At the Application Server command line, run `/opt/netiq/cloudmanager/configurator/config`.
3. Select `NetIQ Cloud Manager - Upgrade Resource Pools` and deselect all other options.
4. Follow the prompts in the utility until the script is complete.

If any zones were not upgraded in the Orchestration Server correctly or if other problems occur, the configurator utility notifies you of the problem.

21 Orchestration Components Upgrade Overview

With this release, upgrade logic you implement for existing installations uses internal version numbers that might seem incongruent. The following Cloud Manager Orchestration 3.4 components are upgraded to Cloud Manager Orchestration 3.5 components:

- Cloud Manager Orchestration Server
- Cloud Manager Orchestration Console
- Cloud Manager Orchestration Agent
- Cloud Manager Monitoring Server
- Cloud Manager Monitoring Agent

This section explains what you can expect from a Cloud Manager Orchestration 3.4 components upgrade to Cloud Manager Orchestration 3.5 components.

- Section 21.1, “Basic Functions of the Orchestration Components Upgrade,” on page 187
- Section 21.2, “Cloud Manager Orchestration Components That Are Not Upgraded,” on page 188

21.1 Basic Functions of the Orchestration Components Upgrade

Before you begin the Orchestration components upgrade process, you need to know the underlying assumptions of the process so that you can better understand how to proceed. The following list details the most important of those assumptions:

- To check the installed Cloud Manager Orchestration components for version number:

  Linux: Run the following command on a Linux machine where agent, client, or server components are installed.

  ```
  rpm -qa | grep 'novell-zen'
  ```

  Windows: On a Windows machine, open the **Add or Remove Programs** console in Windows and look for the agent or client version number in the programs list.
- The upgrade to Cloud Manager Orchestration 3.5 must be done for all Orchestration Servers, Orchestration Consoles, and all Orchestration Agent components. Running older agents with newer server components or running older Orchestration Consoles and interfaces with newer server components (or vice versa) is not supported.
- Upgrading a prior release of a 32-bit Cloud Manager Orchestration Server installation to a newer 64-bit version of Cloud Manager Orchestration Server is not supported. Similarly, upgrading a prior release of a 64-bit Cloud Manager Orchestration Server installation to a newer 32-bit version of Cloud Manager Orchestration Server is not supported.
- The Orchestration Server must be upgraded before Orchestration Agents are upgraded. The Cloud Manager Orchestration Server operates with older agents running, but newer 3.5 agents cannot communicate with the Cloud Manager Orchestration Server 3.4. You can upgrade the

agents by selecting the Upgrade option on the Resource Registration dialog in the Cloud Manager 3.5 Orchestration Console. For more information, see Chapter 10, “Creating a Resource Account,” on page 73.

- After you upgrade the server components, older versions of the Orchestration Agents, the Orchestration Console might not work with the newer server components. The Cloud Manager Orchestration Console identifies the managed nodes that have non-compatible agents. For more information about component compatibility, see Appendix B, “Compatibility Checking Behavior for Orchestration Components,” on page 207.
- After an upgrade to Cloud Manager Orchestration Server 3.5, some specific provisioning adapter jobs (such as vsphere) and existing VMs will not be available for use, even though their files still reside on the Orchestration Server. These provisioning adapter jobs need to be reconfigured.
- If errors occur during the upgrade process, you can attempt to resolve those errors and run the upgrade process again. For more information about how this recovery works, see Section C.1, “Upgrade Failure Scenarios,” on page 211.
- After the Cloud Manager Orchestration Server 3.4 is upgraded to Cloud Manager Orchestration Server 3.5, rolling back to Cloud Manager Orchestration Server 3.4 is not supported.
- Step-by-step information about the events occurring during the upgrade process is recorded in `server.log`, located in the `/var/opt/novell/zenworks/zos/server/logs` directory.

  In some situations, the server log might not exist. You can also check the install log at `/var/opt/novell/novell_zenworks_orch_install.log` for upgrade information.

If you understand what to expect from the upgrade, you are ready to proceed to Chapter 22, “Upgrading Orchestration Server Components,” on page 189.

21.2 Cloud Manager Orchestration Components That Are Not Upgraded

When you upgrade from Cloud Manager Orchestration 3.4 to Cloud Manager Orchestration 3.5, the core components are not upgraded or redeployed. Instead, the old core components are replaced with new core Orchestration components. If you made any changes to the original core components, those changes are saved, so you can manually re-enter the custom configuration you want after the upgrade.

NOTE: After the upgrade to Cloud Manager Orchestration 3.5 components, some earlier-version provisioning adapter jobs (vsphere) and the VMs provisioned by those jobs are not redeployed for use. Any resource or other objects previously managed by these provisioning adapters are no longer manageable, even though they still exist in the Cloud Manager Orchestration Server.

---

22 Upgrading Orchestration Server Components

This section provides information about upgrading from earlier Cloud Manager Orchestration components to current Cloud Manager Orchestration components. It is important that you upgrade the Orchestration components you have installed in the sequence that follows.

- Section 22.1, “Backing Up Orchestration Components Prior to Upgrading,” on page 189
- Section 22.2, “Checking the Current Version of Cloud Manager Orchestration Components,” on page 190
- Section 22.3, “Snapshotting the Existing Orchestration Server Installation,” on page 190
- Section 22.4, “Upgrading Orchestration Components,” on page 191
- Section 22.5, “Configuring the Upgraded Orchestration Components,” on page 195
- Section 22.6, “Configuring the Remote Audit Database after Orchestration Components Are Upgraded,” on page 199
- Section 22.7, “Running Discovery on VM Hosts and Images,” on page 200
- Section 22.8, “Alternate Methods for Upgrading Older Agents and Clients,” on page 200
- Section 22.9, “Running the Upgrade Configuration on an Enterprise Scale,” on page 202
- Section 22.10, “Upgrading a Orchestration High Availability Configuration,” on page 203

NOTE: To perform a mass upgrade of Cloud Manager Orchestration components, we recommend that you use a reputable application software distribution method to upgrade to the current Orchestration version shipping with Cloud Manager.

For more information, see Section 22.9, “Running the Upgrade Configuration on an Enterprise Scale,” on page 202.

22.1 Backing Up Orchestration Components Prior to Upgrading

As with the installation of any software, it is always a wise precaution to back up a working copy of Cloud Manager components before you install the newer version of Cloud Manager Orchestration components.

To back up the old Orchestration components:

1. Make a copy of the directories under the `/var/opt/novell/zenworks/zos/` directory.
2. Back up the `/opt/novell/zenworks/zos/server/license/key.txt` file.

When you want to use the older version of Cloud Manager Orchestration, stop any instance of the Orchestration Server or the Cloud Manager Monitoring server, copy the backup directory you made earlier to its original location, then start the server from this location. Follow this same procedure for the Orchestrate Agents.

22.2 Checking the Current Version of Cloud Manager Orchestration Components

Before you upgrade the Cloud Manager Orchestration packages from version 3.4 to the Cloud Manager Orchestration 3.5 packages, you should check which packages need to be upgraded and which non-Orchestration packages are included in the product packages.

To do this, run the following command:

```
rpm -qa | grep 'novell-zen'
```

We recommend that you record the results of this command so that you can compare it with the results of a similar task following the upgrade. See Section 22.4.4, “Checking the Upgraded Version of the Orchestration Components,” on page 194.

Orchestration Agents must be the same version as the Orchestration Server in order to facilitate full functionality of the product. When you upgrade an Orchestration Server, you must upgrade its agents.

22.3 Snapshotting the Existing Orchestration Server Installation

Before you begin the upgrade process of the Orchestration Server, ensure that all running jobs are complete and take a snapshot of the current configuration of the Orchestration server. If the jobs have not completed on their own, the upgrade processes forcibly cancels them, which is the normal behavior when the server is shut down. The effect on the jobs is that they are terminated abruptly before they finish running. The specific consequence of this termination depends on the job that is terminated.

When you are sure that the jobs are complete, you need to run a specific shutdown command to prepare a snapshot of the current configuration of the server so that a new version of a server can be started with the configuration of the old server.

When an upgrade of server components occurs, all of the current server settings (configuration) and state (model) for the current instance is written to a platform-independent XML encoded snapshot. This snapshot is read in by a newly upgraded server instance to initialize its settings and state to that of the previous server instance.

The snapshot data is read when a newly upgraded server instance is first started, initializing its settings and its state to that of the previous server instance. The snapshot files must exist in `/var/opt/novell/zenworks/zos/server/snapshot`.

To perform the snapshot:

1. Check the running status of the server:

   ```
   /etc/init.d/netiq-cmosserver status
   ```

   If the Orchestration Server is already stopped, you must start it before a snapshot can be created:

   ```
   /etc/init.d/netiq-cmosserver start
   ```

2. Create a snapshot of the server’s current configuration with the following command:

   ```
   /etc/init.d/netiq-cmosserver stop --snapshot
   ```

You can also create the snapshot by using the Orchestration Console to shut down the server. To do so, select **Server** > **Shutdown Server** to display the Server Shutdown Confirmation dialog.



Select **Perform Snapshot of Server State**, then click **Shutdown**.

22.4 Upgrading Orchestration Components

There are two methods for upgrading Orchestration Server packages.

- If you want to use a graphical user interface (GUI) see “Upgrading Orchestration Packages Using YaST2” on page 191.
- If you want to use the command line to upgrade, see “Upgrading Orchestration Packages Using the zypper Command” on page 193.

22.4.1 Upgrading Orchestration Packages Using YaST2

Use the following procedure if you want to use YaST, a graphical user interface, to upgrade the Cloud Manager Orchestration packages.

1 Download the Cloud Manager 2.5 ISO (64-bit), then prepare it for installation:

TIP: The Cloud Manager 2.5 product ISO includes the packages for Cloud Manager Orchestration components, which carry the 3.5 version.

1a (Optional) Burn a DVD of the ISO image and load it into the DVD drive of the target machine.

1b (Optional) Copy the ISO image to the local file system.

To mount the ISO image file on a particular machine,

1b1 Log in to the target server as `root`.

1b2 Open YaST2.

1b3 In the YaST Control Center, click **Software**, then click **Software Repositories** to display the Configured Software Repositories view.

1b4 In the Configured Software Repositories view, click **Add** to open the Media Type view.

1b5 In the Media Type view, select **Local ISO Image**, then click **Next** to open the Local ISO Image view.

1b6 In the **Repository Name** field, enter a name for the repository.

1b7 In the **Path to ISO Image** field of the Local Directory or ISO view, browse to the path where you copied the ISO image file, then click **Next**.

1c (Optional) Mount the ISO image file on the machine where Cloud Manager Orchestration is to be installed (the “target” machine).

If you want to mount the ISO image file on a particular machine,

1c1 Log in to the target server as `root`.

1c2 From the command line of the target machine, enter the following commands

```
mkdir /mnt/iso

mount -o loop NetIO_Cloud_Manager-2.5<SLES_version>.x86_64.iso /mnt/iso
```

(where you substitute the name of the ISO (64-bit) that you are using).

1c3 Open YaST2.

1c4 In the YaST Control Center, click **Software**, then click **Installation Source** to display the Configured Software Catalogs view.

1c5 In the Configured Software Catalogs view, click **Add** to open the Media Type view.

1c6 In the Media Type view, select **Local Directory**, then click **Next** to open the Local Directory view.

1c7 In the **Repository Name** field, enter a name for the repository.

1c8 In the **Path to Directory** field of the Local Directory view, enter the mount point:

```
/mnt/iso
```

1d (Optional) If you are installing the ISO image to a large network, extract the product files from the ISO image to a Web server / FTP server that can be accessed by the target machine without the need for authentication or anonymous login.

To add an `.iso` file or Web URL as an installation source in YaST,

1d1 Log in to the target SLES server as root, then open YaST2.

1d2 In the YaST Control Center, click **Software**, then click **Installation Source** to display the Configured Software Catalogs view.

1d3 In the Configured Software Catalogs view, then click **Add** to open the Media Type view.

1d4 In the Media Type view, select an installation media type.

1d4a (Example) If you extracted the ISO image to a Web Server or FTP Server, select **HTTP** (or **FTP**), then click **Next** to open the Server and Directory view.

1d4b In the **Server Name** field of the Server and Directory view, enter the Server Name (IP Address or DNS Name), in the **Directory on Server Field**, enter the directory name where you extracted the ISO, then click **Next**.

2 Upgrade Orchestration Server software packages to Orchestration Server software packages:

2a Log in to the target SLES server as `root`, then open YaST2.

2b In YaST2, select **Software** > **Software Management**, select the method to open the product ISO on your machine, click **Next**, then follow the procedures to mount the ISO.

2c From the License Agreement page, select the option to agree to the license terms, then click **Next**.

2d In YaST2, open the **Filter** drop-down list, select **Patterns** or **Install Sources** to display the Patterns and Packages view, then click **Details** to close the information pane and open the Package frame.

2e In the **Patterns** frame (left-hand side of the view), select a Cloud Manager Orchestration pattern already installed on this server. The **Package** frame lists the packages either installed or not yet installed for this pattern.

Component packages already installed to the server are checked.

NOTE: Package names for this release of Cloud Manager Orchestration continue to use “novell-zenworks” in the prefix or “Cloud Manager Orchestration” in the summary description.

2f Right-click on any of the installed package names, click **All in This List** > **Update if newer version available**.

2g Add the new **novell-zenworks-zos-server-data-livecd** package, then click **Accept** to install the upgraded packages.

2h Repeat Step 2e through Step 2g for each installed pattern you are upgrading.

After the RPM files are upgraded, scripts are run that do the following:

- Back up the existing server instance directory
- Upgrade the RPM files for the selected Orchestration patterns

3 Configure the Cloud Manager Orchestration Server. You can use one of two information gathering methods to perform the configuration:

- Run the Orchestration Server product configuration script. If you use this method, continue with the steps in “Configuring the Upgraded Orchestration Components” on page 195.
- Run the GUI Configuration Wizard. If you use this method, skip to the steps in “Configuring the Upgraded Orchestration Components” on page 195.

22.4.2 Upgrading Orchestration Packages Using the zypper Command

Use the following procedure if you want to use `zypper` commands to upgrade the Cloud Manager Orchestration packages on SLES machines. If you want to use the GUI Configuration Wizard to upgrade, see “Configuring an Upgraded Orchestration Agent Installed on the Server” on page 198.

For more zypper commands, see “Other Useful zypper Commands for Upgrade” on page 194.

1 Download the appropriate Cloud Manager 2.5 ISO, then prepare it for installation:

- (Optional) Burn a DVD of the ISO image, mount the DVD, then extract the contents of the `.iso` folder to the local file system of the server.
- (Optional) Extract the contents of the `.iso` folder to the local file system of the server.

2 At the command line, change to the directory where the Cloud Manager `.iso` folder was extracted, then run the commands to upgrade Cloud Manager 2.4 to Cloud Manager 2.5:

2a Run the following command:

```
zypper sa -t yast2 "http://<ip_address_of_local_server>/
<directory_location_of_iso_files>"
```

Alternative 1: If you have chosen not to extract the files and you want to use the `.iso` image to upgrade, use the following command:

```
zypper sa -t yast2 "iso:/?iso=<directory_location_of_iso>/<iso_name>"
<service_name>
```

or

```
zypper sa -t yast2 "iso:/?iso=<directory_location_of_iso>/<iso_name>"
"<repo_alias>"
```

For example, for the ISO located at `/root/Desktop/NetIQ_Cloud_Manager-2.5-SLE11.x86_64.iso`, you could use this command:

```
zypper sa -t yast2 "iso:/?iso=/root/Desktop/NetIQ_Cloud_Manager-2.5-
SLE11.x86_64.iso" "CMOS_Server"
```

Alternative 2: If you are using an ftp server and you want to use the `.iso` image to upgrade, use the following command:

zypper sa -t `yast2 "ftp://<ip_address_of_local _server>/<directory_location_of_iso_files>"`

2b Run the following command:

```
zypper ref <repo_alias>
```

2c Run the following command:

```
zypper dup -r <repo_alias>
```

22.4.3 Other Useful zypper Commands for Upgrade

You might find the other zypper commands listed in the table below to be useful during the server upgrade process.

Table 22-1 *zypper Commands That Might Be Useful During Server Upgrade*

Command	Description
`zypper refresh $REPO_ALIAS`	Builds metadata and cache.
`zypper pa $REPO_ALIAS`	Displays all packages in the repository.

22.4.4 Checking the Upgraded Version of the Orchestration Components

After you upgrade the Cloud Manager Orchestration packages to Cloud Manager Orchestration 3.4 components, you should check the upgraded software packages to confirm that all of the earlier versions of the product components are now updated and which of the non-NetIQ packages have been updated.

To do this, change to the directory where the current version of Cloud Manager Orchestration components were extracted, then run the following command:

```
rpm -qa | grep 'novell-zen'
```

Compare the results of this command with the results you had with the check you performed before the upgrade (see Section 22.2, “Checking the Current Version of Cloud Manager Orchestration Components,” on page 190). If some of the components have not been upgraded from the earlier version, the incompatibility between the components could cause unexpected behavior.

22.5 Configuring the Upgraded Orchestration Components

This section discusses the basic upgrade configuration of all Cloud Manager Orchestration components after each is upgraded. Component configuration is done either with a text-based configuration tool or with a GUI Wizard configuration tool.

The text-based configuration script detects which RPM patterns are installed, but the GUI Configuration Wizard requires that you specify the components to be configured, whether the patterns have been installed on the server or not.

Both the text-based tool and the GUI Wizard tool produce a configuration file (`/etc/opt/novell/novell_zenworks_orch_install.conf`).

The section includes the following information:

- Section 22.5.1, “Some Considerations When Configuring with the GUI Wizard,” on page 195
- Section 22.5.2, “Configuring the Upgraded Orchestration Server,” on page 195
- Section 22.5.3, “Configuring the Upgraded Monitoring Server,” on page 198
- Section 22.5.4, “Configuring an Upgraded Orchestration Agent Installed on the Server,” on page 198

22.5.1 Some Considerations When Configuring with the GUI Wizard

If you have only a keyboard to navigate through the pages of the GUI Configuration Wizard, use the Tab key to shift the focus to a control you want to use (for example, a **Next** button), then press the Spacebar to activate this control.

When you have finished answering the configuration questions in the wizard, the Cloud Manager Orchestration Configuration Summary page displays. Although this page of the wizard lets you navigate by using the Tab key and the Spacebar, you need to use the Ctrl+Tab combination to navigate past the summary list. Click **Back** if you accidentally enter the summary list, and re-enter the page to navigate to the control buttons.

By default, the **Configure now** check box on the page is selected. If you accept this default, the wizard starts the Orchestration Server and applies the configuration settings. If you deselect the check box, the wizard writes out the configuration file to `/etc/opt/novell/novell_zenworks_orch_install.conf` without starting the Orchestration Server or applying the configuration settings. For more information, see “Using the GUI Configuration Wizard to Run a Delayed Configuration of the Orchestration Server” on page 198.

22.5.2 Configuring the Upgraded Orchestration Server

Because so much of Cloud Manager’s operations depends on the Orchestration Server, we recommend that you configure it before you configure any other Orchestration component. Use either of the following methods to configure your upgraded Orchestration Server.

- “Running the Cloud Manager Orchestration Configuration Utility” on page 196
- “Using the GUI Configuration Wizard to Run a Delayed Configuration of the Orchestration Server” on page 198

Running the Cloud Manager Orchestration Configuration Utility

1 Ensure that you are ready with the information that you will be prompted for during the configuration procedure (GUI or text-based):

Server Upgrade Configuration Requirement	Explanation and Action
Configuration Selection	This section discusses upgrade, so specify `u` (for `upgrade`) in the configuration script, or select the **Upgrade** check box in the wizard.
Configuration Type	Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation. This section discusses standard installation, so specify `s` (for `standard`) or press Enter to accept the default. For more information about High Availability configuration, see Part IV, “Advanced Installation and Integration Topics,” on page 133.
Administrator User	**IMPORTANT:** Do not change the administrator name from the name you used in the original installation.
Administrator Password	Enter the password of the Cloud Manager Orchestration Server administrator you used in the previous installation (for Cloud Manager 2.4/Orchestration Server 3.4). **IMPORTANT:** Do not change the administrator password from the password you used in the original installation.
Confirm Password / Retype Password	Re-enter the administrator password.
Automatic Agent Upgrade	If there are existing Orchestration Agents configured to use this server, they will normally be placed on a pending upgrade list to be approved for automatic upgrade by the administrator. You can enable automatic upgrade of old agents to avoid this approval step if you enter `yes` or check the **Automatic Agent Upgrade** check box (wizard). If you choose to automatically upgrade the agents later, you can do so in the Orchestration Console. For more information, see “Resources Panel” in the “Orchestration Server Object” section of the *Cloud Manager Administrator Reference*.

Server Upgrade Configuration Requirement	Explanation and Action
Upgrade Auditing Database	If you previously enabled auditing, you need to upgrade the database schema. If you use a PostgreSQL database, you can use the configuration utility to upgrade it. If you use a different RDBMS, you need configure it separately. If you select the **Upgrade Audit Database** check box (wizard), or if you enter `yes` for this prompt, you must specify ◆ the JDBC URL you used previously to connect to the audit database. Do not include a database name after the trailing forward slash (/). ◆ the database name you used previously to create the audit database. ◆ the user name for the PostgreSQL audit database your created previously for logging in. ◆ the password for the PostgreSQL audit database you created previously for logging in. You are required to verify this password. If you want to manually configure the database after upgrade, see Section 22.6, “Configuring the Remote Audit Database after Orchestration Components Are Upgraded,” on page 199.
Admin Info Port	You need to verify the port you previously configured for access to the Administration Information page. Ensure that you specify the port used previously, not necessarily the default.
Orchestration Agent Port	You need to verify the port you previously configured for communication between the Orchestration Server and the Orchestration Agent. Ensure that you specify the port used previously, not necessarily the default.
Path to License File	As you upgrade, you need a license key (90-day evaluation license or a full license) to use this product. Depending on your arrangements with NetIQ, you either received a new key or you were given permission to reuse your old key. Specify the path to the network location of either your existing key or the new key that was issued to you.

2 At the computer where you installed the upgraded Cloud Manager Orchestration Server pattern, run the Cloud Manager Orchestration configuration utility of your choice:

```
/opt/novell/zenworks/orch/bin/config
```

or

```
/opt/novell/zenworks/orch/bin/guiconfig
```

3 Follow the prompts to complete the configuration.

Using the GUI Configuration Wizard to Run a Delayed Configuration of the Orchestration Server

If you want to delay the upgrade of an Orchestration Server, you can capture the configuration parameters in the Cloud Manager Orchestration GUI Configuration Wizard and apply them at your discretion.

To run the delayed configuration:

1. Run the script for the Orchestrate Configuration Wizard as follows:

   ```
   /opt/netiq/ncm/orch/bin/guiconfig
   ```

2. Configure the parameters of the server upgrade as described in “Configuring the Upgraded Orchestration Server” on page 195.
3. When you are ready to commit the configuration, deselect the **Configure now** check box so that the wizard can write the configuration file to `/etc/opt/netiq/netiq_ncm_orch_install.conf` without starting Orchestrate or applying the configuration settings.

   NOTE: You can use this `.conf` file to start the Orchestrate Agent and apply the settings either manually or with an installation script. Use the following command to run the configuration:

   ```
   /opt/netiq/ncm/orch/bin/config -rs
   ```

4. Click **Next** to display a message asking whether you want to overwrite the `.conf` response file.
5. To upgrade, you need to overwrite the existing file. When prompted, click **Yes** to overwrite the file and display the configuration page.
6. Click **Next** to begin the upgrade configuration for the Cloud Manager Orchestration Server 3.4 to the Orchestration Server 3.5.

22.5.3 Configuring the Upgraded Monitoring Server

The configuration does not require any input from you as you upgrade the Cloud Manager Monitoring Server.

1. At the computer where you installed the Cloud Manager Monitoring Server pattern, run the configuration utility of your choice:

   ```
   /opt/novell/zenworks/orch/bin/config
   ```

   or

   ```
   /opt/novell/zenworks/orch/bin/guiconfig
   ```

2. Select the Monitoring Server as the component that you want to upgrade.
3. Follow the prompts to complete the configuration of the Monitoring Server.

22.5.4 Configuring an Upgraded Orchestration Agent Installed on the Server

The configuration does not require any input from you if you are upgrading Cloud Manager Orchestration Agent on the same machine where the Orchestration Server is installed.

1. At the computer where you installed the Cloud Manager Orchestration Agent pattern, run the configuration utility of your choice:

`/opt/novell/zenworks/orch/bin/config`

or

`/opt/novell/zenworks/orch/bin/guiconfig`

2 Select the Orchestration Agent as the component that you want to upgrade.

3 Follow the prompts to complete the configuration of the Orchestration Agent.

22.6 Configuring the Remote Audit Database after Orchestration Components Are Upgraded

After you have upgraded the Orchestration Server, you can manually configure the existing audit database using the `audit_db_upgrade.sql` script. This script creates an `actions` table in the database. Use the following procedure to manually upgrade the audit database:

1 On the Orchestration Server host machine, use your favorite editor to edit the script `/opt/novell/zenworks/zos/server/conf/audit_db_upgrade.sql`.

   1a Replace the `${DB_NAME}` variable with the PostgreSQL database name (for example, `zos_db`).

   1b Replace the `${DB_USER}` variable with the PostgreSQL schema owner name (for example, `zos`).

2 Use the following commands to run the modified script as the PostgreSQL database administrator for the remote database:

```
su - postgres
psql -h <psql-server-addr> -d postgres -U postgres -f audit_db_upgrade.sql
```

3 Use the following command to log into PostgreSQL, using the database name and schema owner substituted in Step 1 above:

```
su - postgres
psql -h <psql-server-addr> -d zos_db -U zos -f audit_db_upgrade.sql
```

4 Confirm that the database user name and password match the values used when creating the schema owner database user in Section 7.1, “Configuring the Orchestration Server,” on page 60. In this example, the user name is `zos` and the password is `zos`.

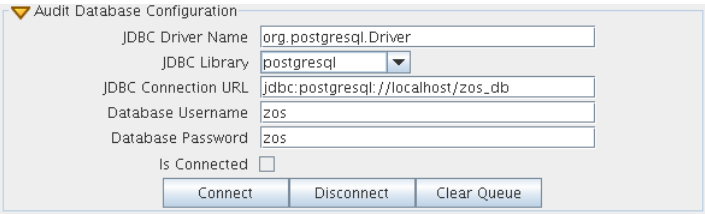

5 Confirm that the database user name and password match the values you replaced in the variables of the `.sql` script. In this example, the user name is `zos` and the password is `zos`.

6 Click **Connect**.

The **Is Connected** check box is selected: the Orchestration Server is connected to the database so that any queued data and subsequent job, user, and resource events are written there.

22.7 Running Discovery on VM Hosts and Images

After the Orchestration components are upgraded and configured, you need to re-discover all of the VMs in the grid so that the new facts are added to the VMs.

To do this from the Orchestration Console **Tools** menu:

1. Click **Provision** > **Discover VM Hosts and Repositories**, select the provisioning adapter you want to run for the discovery, then click **OK**.
2. Click **Provision** > **Discover VM Images**, select the provisioning adapter you want to run for the discovery, then click **OK**.

22.8 Alternate Methods for Upgrading Older Agents and Clients

It is likely that you have installed older-versioned Agents and Clients on machines other than where the Cloud Manager Orchestration 3.4 components were installed. This section includes information that helps you to walk through the upgrade of those agents and clients.

- Section 22.8.1, “Automatically Upgrading the Orchestration Agent from the Cloud Manager Orchestration Console,” on page 200
- Section 22.8.2, “Using the ISO to Upgrade the Orchestration Agent on Red Hat Enterprise Linux 6 Machines,” on page 201
- Section 22.8.3, “Using the ISO to Upgrade the Old Orchestration Agent or the Orchestration Clients on Windows Machines,” on page 202
- Section 22.8.4, “Using the Administrator Information Page to Upgrade the Agents and Clients,” on page 202

NOTE: To perform a mass upgrade of Cloud Manager Orchestration Agents, we recommend that you use a reputable application software distribution method to upgrade to the newer versions that ship with Cloud Manager 2.5.

For more information, see Section 22.9, “Running the Upgrade Configuration on an Enterprise Scale,” on page 202.

22.8.1 Automatically Upgrading the Orchestration Agent from the Cloud Manager Orchestration Console

The Cloud Manager Orchestration Console includes a feature that lets you automatically upgrade older Orchestration Agents on resources (virtual or physical) that connect to the Orchestration Server.

If your grid includes older (that is, older than Cloud Manager version 2.5) Resource objects, their Orchestration Agents cannot connect to a recently upgraded Orchestration Server. This is shown when the Resource Registration icon in the Resource Monitor of the Orchestration Console displays a “flag up”  status.

When you click the Resource Registration icon, the Resource Registration Monitor dialog displays all of the older resource objects attempting to connect. The **Upgrade** option is also available in the dialog. You can select this option, along with all of the older agents that you want to upgrade.

When you use this automatic upgrade method, only the agent rpm (that is, the `zw_zos_agent` pattern) is upgraded, but other Orchestration components that might be present on the resource machine are not upgraded. These include the following components (listed by pattern names):

- `zw_zos_clients`
- `zw_mon_agent`
- `zw_vm_builder`

To separately upgrade these components, you need to use the YaST upgrade from the product ISO (see "Upgrading Orchestration Packages Using YaST2" on page 191) or the zypper upgrade method, which is executed from the bash prompt. For more information, see "Upgrading Orchestration Packages Using the zypper Command" on page 193.

22.8.2 Using the ISO to Upgrade the Orchestration Agent on Red Hat Enterprise Linux 6 Machines

Use the following procedure if you want to use the Add-on method to upgrade the Cloud Manager 3.4 Orchestration Console to a Cloud Manager 3.5 Orchestration Console running on a Red Hat Enterprise Linux (RHEL) 6 machine.

1. Shut down the old Orchestration Console on the machine where you intend to install the new Cloud Manager 3.5 Orchestration Console.
2. Download the appropriate Cloud Manager 2.5 ISO to an accessible network location.
3. Mount the Cloud Manager ISO as a loopback device as in the following example:

   ```
   mount -o loop NetIQ_Cloud_Manager-2.5-SLE11.x86_64.iso /mnt
   ```

   This mounts the ISO in the `/mnt` folder.
4. Navigate to the directory path where the RHEL packages reside. For example:

   ```
   cd /mnt/Linux
   ```

   There are four packages in the `/RHEL6` directory:

   ```
   netiq-ncm-cmos-java-1.7.0_sun_update9-3.x86_64.rpm
   novell-zenworks-orch-config-3.5-<build>.noarch.rpm
   novell-zenworks-orch-config-gui-3.5-<build>.noarch.rpm
   novell-zenworks-zos-agent-3.5-<build>.x86_64.rpm
   ```
5. Use the `rpm` command to install the packages:

   ```
   rpm -Uvh *.rpm
   ```

   If you encounter an issue regarding missing dependencies, you can use the `up2date` command to download and install those. For example, if you were missing `libcaplso` or `libcap.so.1`, you would run the following:

   ```
   up2date --solvedeps=libcp.so,libcap.so.1
   ```
6. (Optional) Increase the heap size that the JVM handles to enable the console to manage a large number of objects.

   6a Open the `cmoc` bash shell script at `/opt/netiq/ncm/cmos/server/bin`.

   6b Inside the script, find the following line where the JVM parameters are defined:

   ```
   JVMARGS="-Xmx256m -Xms256m -Xmn64m -XX:NewSize=64m -XX:MaxNewSize=64m"
   ```

   The `-Xmx` argument specifies the maximum heap size for the JVM. Increasing the heap size prevents a JVM out of memory condition.

6c Change the value in the `-Xmx` argument from 256MB to 512MB.

6d Save the script.

22.8.3 Using the ISO to Upgrade the Old Orchestration Agent or the Orchestration Clients on Windows Machines

The Orchestration Agent and the Orchestration Console are supported on the operating systems specified in Chapter 2, “Cloud Manager System Requirements,” on page 17. To upgrade, install the new Cloud Manager 2.5 release of the Orchestration Agent or the Orchestration Console.

IMPORTANT: When upgrading the Cloud Manager Orchestration Console on a Windows machine, you must uninstall the prior version first, then install the new version of Cloud Manager Orchestration Console.

Use the following steps to download the Orchestration component you want to install:

1. Download the appropriate Cloud Manager ISO to an accessible network location.
2. Create a DVD from the ISO or use a tool that will mount the ISO.
3. Navigate to the directory path where the Windows packages reside.
4. Double-click the appropriate file (`.exe`) to launch an installation and configuration wizard for the Orchestration component.

22.8.4 Using the Administrator Information Page to Upgrade the Agents and Clients

The Administrator Information Page includes installers for the Cloud Manager Orchestration Agents and Clients for Windows and various Linux/UNIX machines (see Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 41). The page has no facility for upgrading an agent or client.

To upgrade, we recommend that you use the methods native to the OS to install the new Cloud Manager release of the Orchestration Agent or Orchestration Console.

IMPORTANT: When you upgrade the Cloud Manager Orchestration Console on a Windows machine, you must uninstall the prior version first, then install the new version.

22.9 Running the Upgrade Configuration on an Enterprise Scale

If you have a number of Server or Agent components to upgrade in an enterprise environment, you might want to follow these general steps to accomplish the upgrade.

1. Use a reputable configuration management tool to distribute and install the upgrade software. Examples include ZENworks Linux Management, ZENworks Configuration Management, and the Red Hat Network.
2. Configure the upgraded components on a base machine, then, use the configuration software to distribute the respective `.conf` files to the servers or nodes being upgraded.

22.10 Upgrading a Orchestration High Availability Configuration

This section provides the steps you need to follow to upgrade a prior version of the Cloud Manager environment to a Cloud Manager high availability configuration. Use the following steps for the upgrade:

1. Using the high availability manager (such as Heartbeat 2), shut down the Cloud Manager Orchestration service.
2. Manually bind the cluster IP address using the following command:

   `ip addr add {CIDR_IP_ADDRESS} dev {Ethernet Device}`
3. Manually start a Cloud Manager Orchestration Server instance on the first node in the cluster (this should be a node that does not include high-availability components) using the following command:

   `/etc/init.d/novell-zosserver start`
4. Manually stop the instance with snapshot. See Section 22.3, “Snapshotting the Existing Orchestration Server Installation,” on page 190 for more information.
5. Manually unbind the IP address using the following command:

   `ip addr del {CIDR_IP_ADDRESS} dev {Ethernet Device}`
6. Upgrade the RPM files on the first node (the one you started in Step 3) of the cluster. For more information, see Section 22.4, “Upgrading Orchestration Components,” on page 191.
7. Run the configuration script. For more information, see “Configuring the Upgraded Orchestration Components” on page 195.
8. Manually bind the IP address using the following command:

   `ip addr add {CIDR_IP_ADDRESS} dev {Ethernet Device}`
9. Start the Orchestration Server instance on the upgraded node (see Step 3 and Step 6).
10. Verify the upgraded server state by attempting to connect to the server (with the Orchestration Console, for instance).
11. Manually stop the following services using the `/etc/init.d/<service_name> stop` command:
    - novell-zosserver
    - apache2
    - novell-gmond
    - novell-gmetad

    where `<service_name>` is `novell-zosserver`, `novell-gmond`, `novell-gmetad`, or `apache2`, and `stop` is the action you want to perform.

    IMPORTANT: Do not snapshot to stop the server instance.
12. Manually unbind IP address using the following command:

    `ip addr del {CIDR_IP_ADDRESS} dev {Ethernet Device}`
13. Upgrade the RPM files on the second node of the cluster.
14. Using the high availability manager (see Step 1), start the Orchestration Server instance.

A Restoring Cloud Manager Application Server in the Event of a System Failure

If a failure occurs in your Cloud Manager 2.5 Application Server system, you can revert to Cloud Manager 2.4 Application Server.

To restore Cloud Manager Application to the previous release:

1. Ensure that all Cloud Manager services are shut down. Use the following command at the Application Server command line:

   ```
   /etc/init.d/netiq/cloudmanager clean
   ```

2. Uninstall the Cloud Manager 2.5 RPM files. Use the following command at the Application Server command line:

   ```
   rpm --erase netiq-cloudmanager-2.5 netiq-cloudmanager-lib-2.5
   ```

3. Install the Cloud Manager 2.4 RPM files. Use the following command at the Application Server command line:

   ```
   rpm -Uvh netiq-cloudmanager-2.4.0-<build_number>.noarch.rpm
   ```

4. Restore the configuration file backup that the configurator collected during the upgrade. Use a command similar to the following:

   ```
   ./backup -r -f /var/opt/netiq/cloudmanager/cfg-auto-backup-<old version>-to-<new version>.tar.gz
   ```

   NOTE: If you chose to perform a complete Cloud Manager file system backup, you would use a command similar to the following:

   ```
   ./backup -r -f /var/opt/netiq/cloudmanager/<full_backup_before_2.5>.tar.gz
   ```

   ---

5. Restore the backup of the PostgreSQL database you used before the upgrade.

   IMPORTANT: The database you have been using with Cloud Manager 2.5 must be clean and the database user must exist for the upgrade to succeed. You might have to drop the existing database, then re-create it before you can revert to the prior database.

   When you are sure that the existing database is clean, use a command like this to restore the former database that you backed up:

   ```
   sudo -u postgres psql -d cloudmanager -f /var/opt/netiq/cloudmanager/
   cm_database_backup.dump.out
   ```

   ---

6. Start the restored version of Cloud Manager:

   ```
   /etc/init.d/netiq/cloudmanager start
   ```

B Compatibility Checking Behavior for Orchestration Components

Managed agents (nodes) report version incompatibility in the agent log file. On the server, the attempted connection by an incompatible agent is detected, and the agent is listed on the Cloud Manager Orchestration Console as incompatible and in need of either an upgrade or downgrade to the correct version. Also, an incompatible agent connection attempt causes the node manager on the server to raise a `NEED_UPGRADE` event that can be caught to provide custom handling of agents in need of upgrade.

The Orchestration Console detects incompatibility only in the Orchestration Agent. The information in the following sections details that behavior.

- Section B.1, “If the Orchestration Server Is Not Compatible with the Orchestration Console,” on page 207
- Section B.2, “When an Agent Version Does Not Match the Server Version,” on page 208

B.1 If the Orchestration Server Is Not Compatible with the Orchestration Console

When the Cloud Manager Orchestration Console detects an older version of the Orchestration Server, the console displays an “old” icon overlay over the grid object.

Figure B-1 *Orchestration Console Displaying an “Old” Icon*

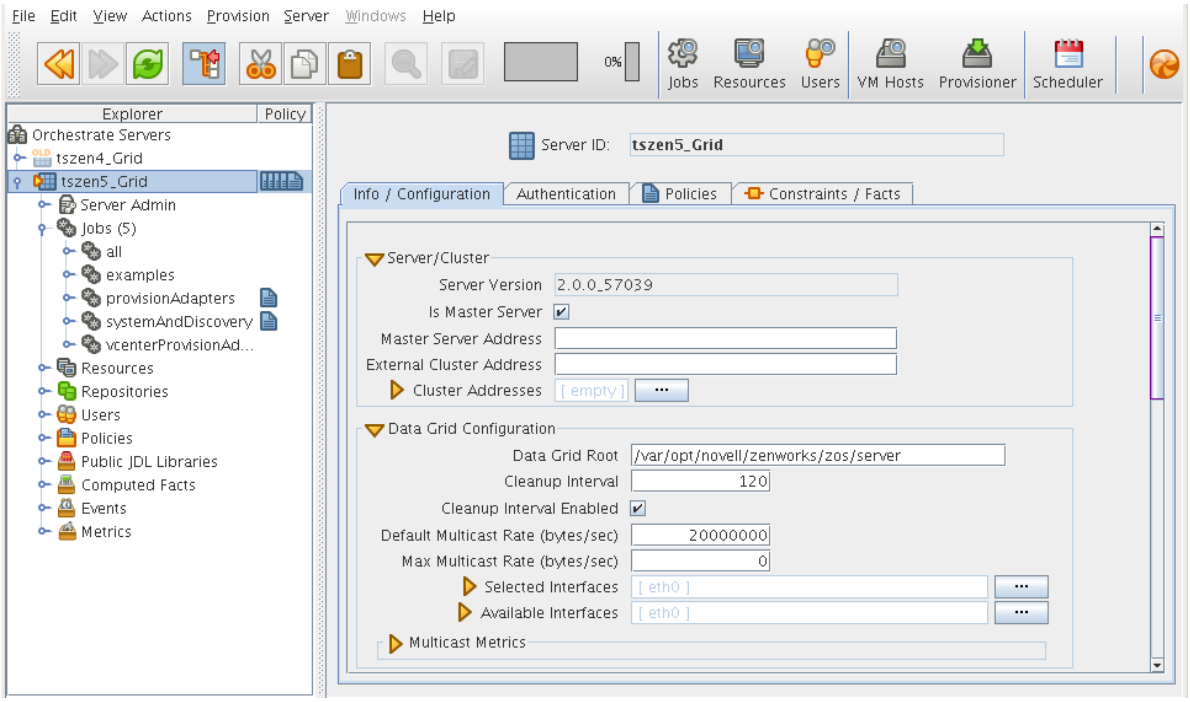

The Orchestration Console displays a “new” icon overlay on the Grid Object if the Grid Object is newer than the console. The version of the server is included in the tool tip display of the grid object in the Explorer tree view. The logged-in server shows the version at the bottom of the view.

B.2 When an Agent Version Does Not Match the Server Version

When an older, incompatible version of the agent communicates with the server, the server detects it and flags the agent as “old.” This incompatibility is displayed in the Orchestration Console, where an older version of the agent is shown in the Tree view with an “old” icon or in the Monitor view with an “old” icon. At this point, the agent also logs a fatal connection error.

Figure B-2 *Old Orchestration Agent Resource Displayed in Tree View*

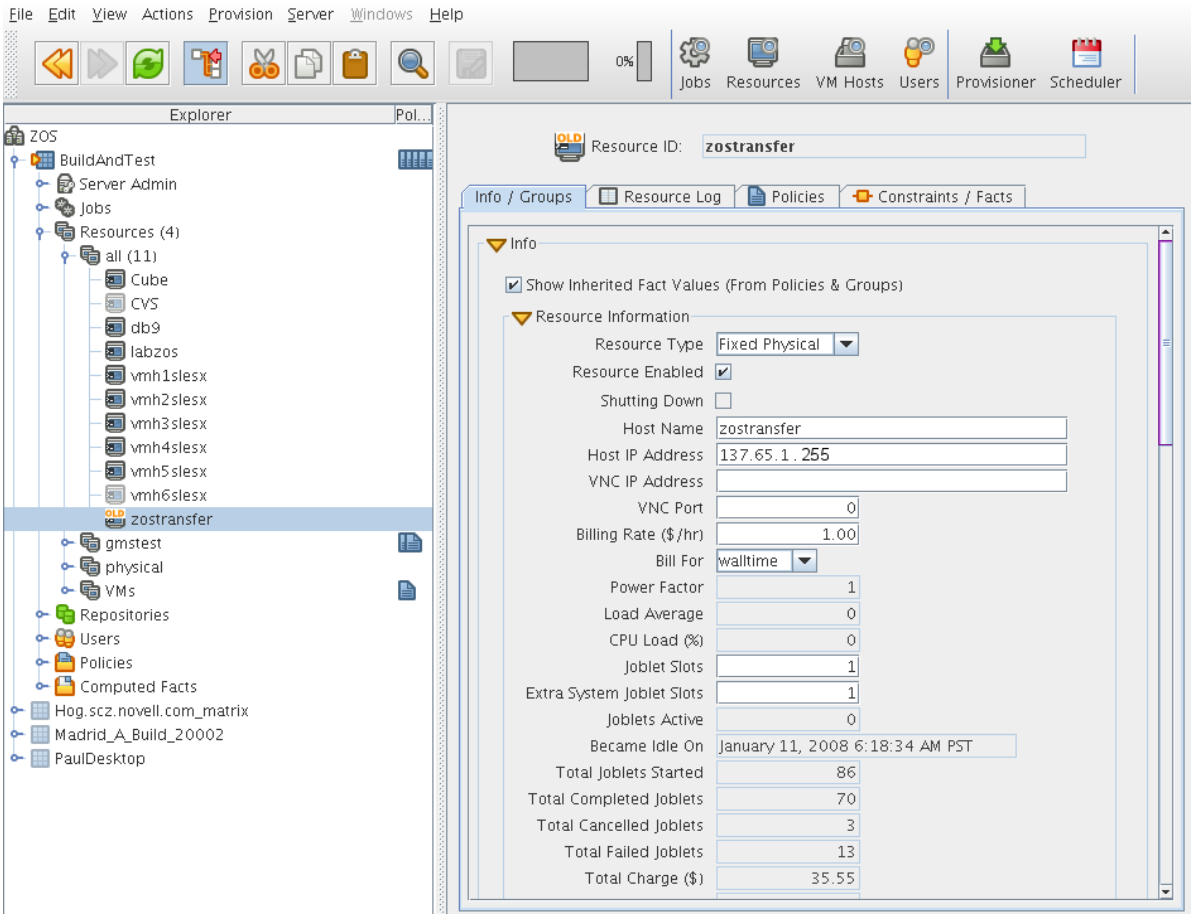

Figure B-3 *Old Orchestration Agent Resource Displayed in Tree View and Monitor View*

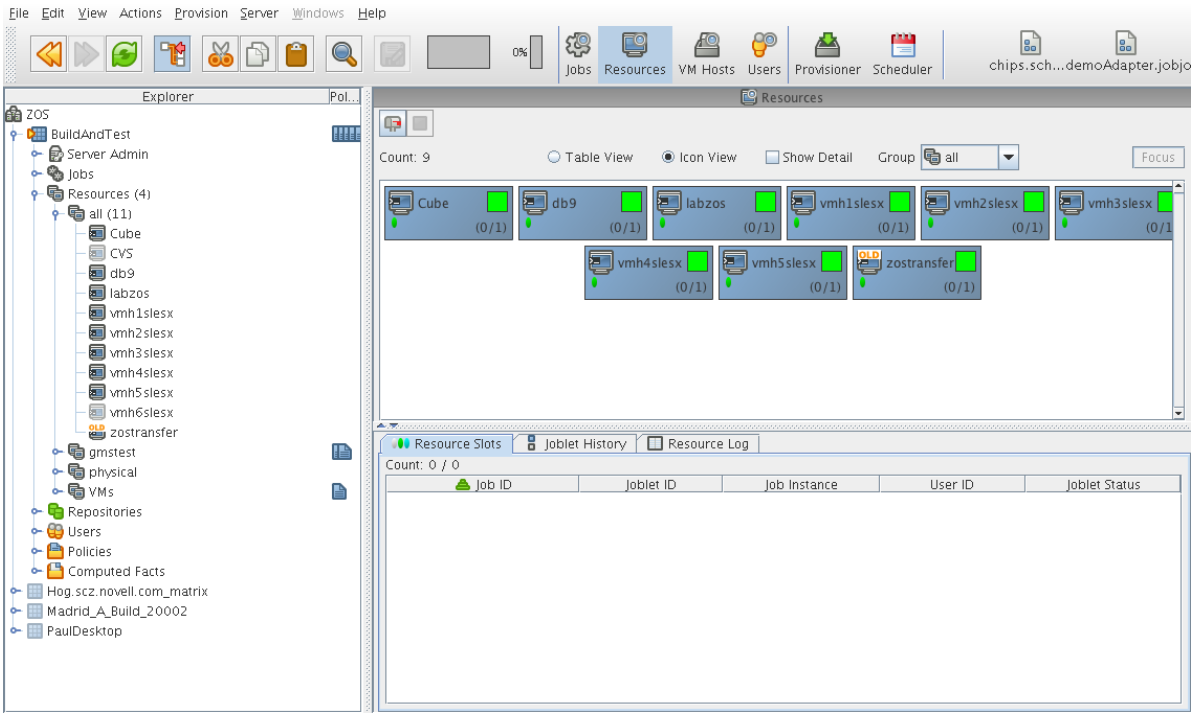

C Recovering from a Failed Orchestration Server Upgrade

The information in this section can help you to recover from a failed Orchestration Server upgrade.

- Section C.1, “Upgrade Failure Scenarios,” on page 211

C.1 Upgrade Failure Scenarios

It is possible that the upgrade process could have problems. If this should occur, we suggest you follow these general steps to recover from those errors and “roll back” to the previous version of Cloud Manager Orchestration.

- Section C.1.1, “Failure Scenario 1: Error Resolution,” on page 211
- Section C.1.2, “Failure Scenario 2: Cannot Resolve Error,” on page 211

C.1.1 Failure Scenario 1: Error Resolution

Follow these steps if you can resolve the error.

1. Open the upgrade log file to learn about the reason for the error, then resolve it.
2. Re-run the configuration

C.1.2 Failure Scenario 2: Cannot Resolve Error

Follow these steps if you cannot resolve the error:

1. Remove the new instance directory for the Cloud Manager Orchestration Server, not including the datagrid.
2. Copy the old instance directory `/var/opt/novell/zenworks/zos.bak` and the license key file to restore the Orchestration Server 3.4 data from the snapshot.
3. Restore the previous version RPM files of the Orchestration Server 3.4 software.

VI Uninstalling Cloud Manager

This section includes information about uninstalling Cloud Manager software:

- Chapter 23, “Uninstalling Orchestration Component Patterns from a SLES Server,” on page 215
- Appendix D, “Documentation Updates,” on page 217

23 Uninstalling Orchestration Component Patterns from a SLES Server

This section provides information about uninstalling Cloud Manager from your data center.

There is no supported wizard or self-contained uninstall tool to remove the Cloud Manager Orchestration Server from your Linux machines, nor is the uninstall feature in YaST2 supported, However, if you no longer want to use the server on a given machine, you can use a Linux package management tool to perform the uninstall. There are a variety of tools you can use to perform the uninstall:

- Use the zypper command line tool to remove the server packages on SLES machines.

  ```
  zypper remove <RPM_Package>
  ```

- Use the YaST setup and configuration tool on SLES machines to select and remove the software packages.
- Use the `rpm -qa` command to look for installed RPM files and the `rpm -e` command to manually remove one or more of them.

To remove the patterns:

1. Shut down all Cloud Manager Orchestration components running anywhere on the grid, including the Orchestration Agent, the Orchestration Server, the Cloud Manager Monitoring Server, and the Cloud Manager Monitoring Agent.
2. Use a package management tool to uninstall the Orchestration Server RPM files:
   - `novell-zenworks-monitor-gmond`
   - `novell-zenworks-zos-agent`
   - `novell-zenworks-zos-server-data-clients`
   - `novell-zenworks-monitor-web`
   - `novell-zenworks-zos-server-data-jre`
   - `novell-zenworks-orch-config-gui`
   - `novell-zenworks-monitor-gmetad`
   - `novell-zenworks-orch-config`
   - `novell-zenworks-zos-clients`
   - `novell-zenworks-zos-server-data-agent`
   - `novell-zenworks-vmwarehouse-base`
   - `novell-zenworks-zos-java`
   - `novell-zenworks-zos-server`
   - `novell-pso-ws`
3. Verify deletion of all of the following directories used by the Orchestration Server:
   - `/opt/novell/zenworks/server`
   - `/var/opt/novell/zenworks/server`
   - `/etc/opt/novell/zenworks/monitor`

- /opt/novell/zenworks/agent
- /var/opt/novell/zenworks/agent
- /root/.novell/zos
- /root/.novell/zoc
- /etc/apache2/conf.d/zos.conf
- /etc/apache2/conf.d/ganglia-auth.conf
- /opt/novell/pso-ws
- /etc/opt/novell/pso-ws
- /var/opt/novell/pso-ws

D Documentation Updates

This section identifies changes made to this guide since the initial release of Cloud Manager 2.5.

- Section D.1, “July 2016,” on page 217

D.1 July 2016

Change	Description
Part V, “Upgrading Cloud Manager,” on page 179	This section was reorganized to clarify that the Application Server should be upgraded before you upgrade Orchestration components.
Chapter 19, “Upgrade Checklist,” on page 181	This section is new.
Appendix A, “Restoring Cloud Manager Application Server in the Event of a System Failure,” on page 205	This section was previously located in the Application Server upgrade section.