

NetIQ Cloud Manager

Installation Guide

May 22, 2013



Legal Notice

NetIQ Product Name is protected by United States Patent No(s): nnnnnnnn, nnnnnnnn, nnnnnnnn.

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

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

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

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

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

© 2013 NetIQ Corporation and its affiliates. All Rights Reserved.

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

Contents

About NetIQ Corporation	9
About this Book and the Library	11
 Part I Preparing for Cloud Manager Installation	 13
1 Installation Checklist	15
 2 Cloud Manager System Requirements	 19
2.1 Cloud Manager Licensing Requirements	19
2.1.1 License Requirements During Installation	19
2.1.2 Product Evaluation Terms and Procurement	19
2.1.3 Upgrading the Server from a Trial License to a Purchased License	20
2.2 Cloud Manager Orchestration Server Requirements	20
2.2.1 Required Network Resources for the Cloud Manager Orchestration Server	21
2.2.2 Required Network Resources for the Cloud Manager Orchestration Web Server	22
2.3 Cloud Manager Orchestration Agent Requirements	23
2.4 Cloud Manager Application Server Requirements	23
2.5 Cloud Manager Application Console Requirements	24
2.6 Requirements and Cloud Manager Support for the Virtual Environment	25
2.6.1 Requirements for Machines Designated as VM Hosts	27
2.6.2 Supported VMs	28
 3 Choosing the Installation Packages and Where to Install Them	 31
3.1 NetIQ Cloud Manager Installation Pattern	31
3.2 Cloud Manager Orchestration Server Install Pattern	32
3.3 Cloud Manager Monitoring Server Pattern	33
3.4 Cloud Manager Orchestration Agent Pattern	34
3.5 Orchestration Console Install Pattern	35
3.6 Monitoring Agent Install Pattern	36
 4 Orchestration Components Preinstallation Tasks	 39
4.1 Gathering Certificate and License Information	39
4.2 Preparing the Server When Multiple NICs and DNS Addresses Exist	39
 Part II Component Installation	 41
 5 Installing Cloud Manager Orchestration Components	 43
5.1 SLES 11 Standard Installation	44
5.2 Alternative Installation Methods for the Orchestration Agent	45
5.2.1 Obtaining the Agent Installer and Supporting Files from the Administrator Information Page	46
5.2.2 Installing the Agent on Windows Machines	47
5.2.3 Manually Installing the Agent Packages on SLES Machines	47
5.2.4 Manually Installing the Agent Linux Packages on RHEL Machines	48

5.2.5	Advanced Agent Installation Methods	49
5.3	Alternative Installation Methods for the Orchestration Console and Clients	52
5.3.1	Obtaining Installers from the Administrator Information Page	52
5.3.2	Installing the Console and Clients on Windows	53
5.3.3	Installing the Console and Clients on a SLES Server	54
5.4	Alternative Installation Methods for the Cloud Manager Monitoring Agent	55
5.4.1	Installing the Cloud Manager Monitoring Agent on Linux Servers	55
5.4.2	Installing the Cloud Manager Monitoring Agent On Windows Machines	56
6	Installing Cloud Manager Application Server Components	59
6.1	Installing to SLES 11	59
Part III	Standard Cloud Manager Component Configuration	61
7	Configuring Cloud Manager Orchestration Components	63
7.1	Configuring the Orchestration Server	64
7.2	Configuring the Monitoring Server and Monitoring Agent	66
7.3	Configuring the Orchestration Agent	66
7.4	Validating and Optimizing the Orchestration Installation	68
8	Configuring Connections to the Cloud Manager Application Server	69
8.1	Enabling a Secure Connection	69
8.1.1	Configuring the Cloud Manager Web Service Secure Port	69
8.2	Enabling a Non-Secure Connection	70
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	75
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 Orchestration Provisioning Adapters	81
11.1	Configuring the vSphere Provisioning Adapter	81
11.1.1	Configuring the vSphere Provisioning Adapter to Discover VMs	81
11.1.2	Discovering Enterprise Resources in Multiple vSphere Environments	92
11.2	Configuring the Citrix XenServer Provisioning Adapter	95
11.2.1	Deploying the Citrix XenServer Provisioning Adapter	96
11.2.2	Configuring the Citrix XenServer Updater	98
11.2.3	Configuring Orchestrator for Personalization with XenServer	98
11.2.4	Using Xen VNC Proxy to Establish a Remote Desktop Connection to XenServer VMs	98
11.3	Configuring the Hyper-V Provisioning Adapter	101
11.3.1	Ensuring that the Orchestration Server Discovers Hyper-V VMs	102

11.3.2	Configuring the Provisioning Adapter to Discover iSCSI Target Repositories	102
11.3.3	Configuring the Provisioning Adapter for Sysprep	102
11.3.4	Enabling a Remote Console Session for a Hyper-V VM	102
11.3.5	Configuring Hyper-V Linux VMs to Enable Visibility of Added vDisks	103
11.4	Configuring the SUSE Xen Provisioning Adapter	103
11.4.1	Cloud Manager Orchestration Defaults in a SUSE Xen Cluster	103
11.5	Configuring the KVM Provisioning Adapter	107
11.5.1	Authentication Settings	107
11.5.2	Debug Settings	108
12	Configuring Sysprep or Autoprep	111
12.1	Understanding and Configuring Sysprep	111
12.1.1	How Sysprep Works	112
12.1.2	Setting Sysprep Facts in the Orchestration Console	112
12.1.3	Using the Sysprep deploy.cab Files	120
12.1.4	Applying Sysprep Facts	122
12.1.5	Example Sysprep Scenarios	123
12.1.6	Known Sysprep Limitations	123
12.2	Understanding and Configuring Autoprep	125
12.2.1	How Autoprep Works	125
12.2.2	Setting Autoprep Facts in the Orchestration Console	126
12.2.3	Applying Autoprep Facts	129
12.2.4	Example Autoprep Scenarios	129
12.2.5	Known Autoprep Limitations	130
13	Using the Cloud Manager Application Server Configuration Tool	131
13.1	Configuring the PostgreSQL Database Connection and Credentials	131
13.2	Configuring the PostgreSQL Database Connection and Credentials	135
13.3	Configuring Cloud Manager to Use Authentication Sources	138
13.3.1	Configuring Authentication to an LDAP Directory	139
13.3.2	Configuring Authentication through an NCSS Director	141
13.3.3	Configuring LDAP Plus NCSS Authentication	143
13.3.4	Configuring Authentication to Novell Access Manager	147
13.4	Installing and Configuring Other Cloud Manager Feature Settings	148
13.4.1	Installing the Cloud Manager Application Console	148
13.4.2	Configuring the Cloud Manager Web Server (Jetty)	148
13.4.3	Configuring the Cloud Manager Web Server to Use SSL	149
13.4.4	Configuring Cloud Manager SMTP Mail Settings	150
13.4.5	Configuring Cloud Manager System Shell Login Information	151
Part IV	Advanced Installation and Integration Topics	153
14	Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support	155
14.1	Overview	155
14.2	Orchestration Server Failover Behaviors	156
14.2.1	Use Case 1: Orchestration Server Failover	156
14.2.2	Use Case 2: VM Builder Behavior at Orchestration Server Failover and Failback	157
14.2.3	Use Case 3: Monitoring Behavior at Orchestration Server Failover and Failback	157
14.3	Installing the Orchestration Server to a SLES 11 Pacemaker Cluster Environment	157
14.3.1	Meeting the Prerequisites	158
14.3.2	Installing the SLES 11 SP2 High Availability Pattern	159
14.3.3	Configuring SLES 11 Nodes with Time Synchronization and Installing Pacemaker to Each Node	159

14.3.4	Setting Up OCFS2 on SLES 11 SP2	161
14.3.5	Installing the Orchestration Server on the First Clustered SLES 11 Node	161
14.4	Configuring the Orchestration Server for High Availability	162
14.4.1	Some Considerations When Configuring with the GUI Wizard	163
14.4.2	The Configuration Procedure	164
14.4.3	Checking the Configuration	165
14.4.4	Running the High Availability Configuration Script	166
14.5	Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster	166
14.6	Creating the Server Cluster Resource Group	167
14.7	Testing the Failover of the Orchestration Server in a High Availability Grid	167
14.8	Installing and Configuring other Orchestration Components to the High Availability Grid	167
14.9	High Availability Best Practices	168
14.9.1	Jobs Using scheduleSweep() Might Need a Start Constraint	168

15 Installing and Configuring the Orchestration Agent for Xen VM Deployment in a SLES HAE Cluster 169

15.1	Xen Cluster Architecture	169
15.2	Installing the Orchestration Agent in a SLES 11 SP1 HAE Xen Cluster	170
15.3	Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster	171
15.3.1	Configuring the Agent for the Cluster	171
15.3.2	Creating the Agent Cluster Resource Group	173
15.3.3	Removing the Orchestration Agent from a Clustered VM Host	174
15.4	Sample Orchestration Agent CIB XML	174

16 Configuring the Orchestration Server to Use an Audit Database 177

16.1	Installing the PostgreSQL Package and Dependencies on an Independent Host	177
16.1.1	Detail	178
16.2	Configuring PostgreSQL to Accept Remote Database Connections	179
16.3	Logging in Locally to the PostgreSQL Database	180
16.4	Creating an Orchestration Server User for the PostgreSQL Database	180
16.5	Configuring the Orchestration Server Audit Database on a Separate Host	180
16.6	Installing and Configuring the Orchestration Server for Use with a Local PostgreSQL Audit Database	182
16.6.1	Installing the PostgreSQL Package and Dependencies	182
16.6.2	Configuring PostgreSQL to Accept Local Database Connections	183
16.6.3	Logging in Locally to the PostgreSQL Database	183
16.6.4	Installing and Configuring the Local Orchestration Server Audit Database	183
16.7	Configuring the Audit Database after the Cloud Manager Orchestration Server Is Configured. . .	185
16.8	Configuring the Remote Audit Database after the Cloud Manager Orchestration Server Is Configured	186
16.9	Modifying Audit Database Tables to Accommodate Long Names	187
16.10	Understanding Grid ID Usage in the Audit Database	187

17 Integrating the Orchestration Server with a Sentinel Collector 189

17.1	Integration Architecture	189
17.2	System Requirements	190
17.3	Importing and Deploying the Orchestration Server Sentinel Collector Plug-in	191
17.4	Connecting the Orchestration Server to the Sentinel Collector Plug-In	193
17.5	Verifying the Sentinel Configuration After Connecting to the Orchestration Server	193
17.6	Event Classification and Taxonomy Keys	194
17.7	Plain Text Visibility of Sensitive Information	197

18 Configuring Secure Authentication Sources to Communicate with Cloud Manager	199
18.1 Configuring Cloud Security Service Connectors for Cloud Manager	199
18.1.1 Creating and Configuring an NCSS Connector for the Cloud Manager Application Console	200
18.1.2 Creating and Configuring an NCSS Connector for the Cloud Manager API	203
18.1.3 Creating and Configuring an NCSS Connector for the Cloud Manager Tasks API	203
18.1.4 Creating and Configuring an NCSS Connector for the Cloud Manager App Services API	203
18.2 Configuring Novell Access Manager to Work with Cloud Manager	204
18.2.1 Managing a Reverse Proxy for Authentication to Cloud Manager	204
 Part V Upgrading	 211
19 Orchestration Components Upgrade Overview	213
19.1 Basic Functions of the Orchestration Components Upgrade	213
19.2 Cloud Manager Orchestration Components That Are Not Upgraded.	214
 20 Upgrading Cloud Manager Orchestration Components	 217
20.1 Upgrading Orchestration Components.	217
20.1.1 Backing Up the Orchestration Components Prior to Upgrading	218
20.1.2 Backing Up the Application Components Prior to Upgrading	218
20.1.3 Checking the Current Version of Cloud Manager Orchestration Components	219
20.1.4 Snapshotting the Existing Orchestration Server Installation	219
20.1.5 Upgrading the Orchestration Packages	220
20.1.6 Checking the Upgraded Version of the Orchestration Components	223
20.1.7 Configuring the Upgraded Packages	224
20.1.8 Manually Configuring the Remote Audit Database after Orchestration Components Are Upgraded	227
20.1.9 Upgrading the XenServer Provisioning Adapter	228
20.1.10 Running Discovery on VM Hosts and Images	228
20.2 Alternate Methods for Upgrading Older Agents and Clients	229
20.2.1 Automatically Upgrading the Orchestration Agent from the Cloud Manager Orchestration Console	229
20.2.2 Using the ISO to Upgrade the Orchestration Agent on Red Hat Enterprise Linux 5 Machines	230
20.2.3 Using the ISO to Upgrade the Old Orchestration Agent or the Orchestration Clients on Windows Machines	231
20.2.4 Using the Administrator Information Page to Upgrade the Agents and Clients	231
20.3 Running the Upgrade Configuration on an Enterprise Scale	231
20.4 Upgrading a Cloud Manager Orchestration High Availability Configuration.	232
 21 Upgrading the Cloud Manager Application Server Components	 233
21.1 Backing Up the PostgreSQL Database	233
21.2 Performing a Complete Cloud Manager System Backup	234
21.3 Running the Cloud Manager Configuration Script	234
21.3.1 Using the Configurator Tool to Update Resource Pool Data on the Cloud Manager Application Server	235
21.4 Upgrading from a Pre-2.1.5 Version of Cloud Manager	235
21.5 Restoring Cloud Manager In the Event of a System Failure	236

A	Compatibility Checking Behavior for Orchestration Components	237
A.1	If the Orchestration Server Is Not Compatible with the Orchestration Console	237
A.2	When an Agent Version Does Not Match the Server Version	238
B	How to Recover from a Failed Orchestration Server Upgrade	241
B.1	Upgrade Failure Scenarios	241
B.1.1	Failure Scenario 1: Error Resolution	241
B.1.2	Failure Scenario 2: Cannot Resolve Error	241
B.2	Restoring the Orchestration Server If the Upgrade Fails	242
B.2.1	Requirements	242
B.2.2	Rollback Procedure Using the rug Command	242
Part VI	Uninstalling	247
22	Uninstalling Orchestration Component Patterns from a SLES Server	249
23	What's Next?	251

About NetIQ Corporation

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

Contacting Sales Support

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

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

Contacting Technical Support

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

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

Contacting Documentation Support

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

Contacting the Online User Community

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

About this Book and the Library

The *Installation Guide* provides information about planning and installing the NetIQ Cloud Manager components. This section includes the following information:

- ♦ [Part I, “Preparing for Cloud Manager Installation,” on page 13](#)
- ♦ [Part II, “Component Installation,” on page 41](#)
- ♦ [Part III, “Standard Cloud Manager Component Configuration,” on page 61](#)
- ♦ [Part IV, “Advanced Installation and Integration Topics,” on page 153](#)
- ♦ [Part V, “Upgrading,” on page 211](#)
- ♦ [Part VI, “Uninstalling,” on page 247](#)
- ♦ [Chapter 23, “What’s Next?,” on page 251](#)

Intended Audience

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

Other Information in the Library

The library provides the following information resources:

Product Overview

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

Installation Guide

Provides detailed planning and installation information.

Procedures Guide

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

Reference Guide

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

Preparing for Cloud Manager Installation

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

- ♦ [Chapter 1, “Installation Checklist,” on page 15](#)
- ♦ [Chapter 2, “Cloud Manager System Requirements,” on page 19](#)
- ♦ [Chapter 3, “Choosing the Installation Packages and Where to Install Them,” on page 31](#)
- ♦ [Chapter 4, “Orchestration Components Preinstallation Tasks,” on page 39](#)

1 Installation Checklist

To ensure that you successfully install and configure NetIQ 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.

Task	Details
<input type="checkbox"/> Review Cloud Manager concepts and terminology	<p>NetIQ 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.</p> <p>If you are not already familiar with Novell Cloud Manager concepts and its interaction with these other products, see the NetIQ Cloud Manager Product Overview.</p>
<input type="checkbox"/> Virtualize your physical datacenter	<p>If you have not already applied a virtualization infrastructure to your physical datacenter, you need to implement a hypervisor technology. NetIQ Cloud Manager supports SUSE Linux Xen, Citrix Xen, VMware vSphere (vCenter), and Microsoft Hyper-V hypervisors and the Amazon EC2 virtual environment.</p>
<input type="checkbox"/> Review the supported Cloud Manager environments and software installation requirements	<p>See Chapter 2, "Cloud Manager System Requirements," on page 19 and Chapter 2.6, "Requirements and Cloud Manager Support for the Virtual Environment," on page 25.</p>
<input type="checkbox"/> Prepare for Orchestration components installation.	<p>Some security certificate and licensing tasks must take place before you begin installing Orchestration components.</p> <p>See Chapter 4, "Orchestration Components Preinstallation Tasks," on page 39.</p>
<input type="checkbox"/> Install the Cloud Manager Orchestration Server, the Orchestration Console, and the Orchestration Agent	<p>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.</p> <p>There are some alternative methods you can use for installing these components. See Chapter 5, "Installing Cloud Manager Orchestration Components," on page 43.</p>

Task	Details
<input type="checkbox"/> Configure the Cloud Manager Orchestration components	<p>Although you could install the Cloud Manager Application components while you have the installation media mounted on the SUSE server, it's not likely that Application and Orchestration components will be installed on the same server. For this reason, you should configure the packages you have installed for the Orchestration components. These configuration tasks include:</p> <ul style="list-style-type: none"> ♦ 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 <p>See Chapter 7, “Configuring Cloud Manager Orchestration Components,” on page 63, Chapter 10, “Creating a Resource Account,” on page 73, Chapter 11, “Configuring Orchestration Provisioning Adapters,” on page 81, and Chapter 8, “Configuring Connections to the Cloud Manager Application Server,” on page 69.</p> <p>NOTE: Although you might not need some of the advanced functionality in Orchestration components, you might be interested in the advanced configuration tasks detailed in Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 45 and in Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 52.</p>
<input type="checkbox"/> Prepare for Cloud Manager installation	<p>Before installing and configuring Cloud Manager, you need to prepare a remote database for storing Cloud Manager data. See Chapter 13.1, “Configuring the PostgreSQL Database Connection and Credentials,” on page 131.</p> <p>You also need to decide which method you want to use to authenticate your users in the Cloud Manager system. The authentication method you use depends on the external authentication source or sources you have already implemented in your data center environment. The sources supported by NetIQ Cloud Manager include:</p> <ul style="list-style-type: none"> ♦ LDAP (Active Directory or eDirectory) ♦ Novell Cloud Security Services (NCSS) ♦ LDAP and NCSS combined ♦ Novell Access Manager (NAM) <p>There are required tasks you need to perform and information you need to gather for to prepare your chosen authentication source to support Cloud Manager authentication and configuration.</p> <p>See Chapter 13, “Using the Cloud Manager Application Server Configuration Tool,” on page 131.</p>

Task	Details
<input type="checkbox"/> Install the Cloud Manager Application Server and its console.	<p>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.</p> <p>See Chapter 6, “Installing Cloud Manager Application Server Components,” on page 59.</p>
<input type="checkbox"/> Configure the Cloud Manager system	<p>After installation, you must complete several configuration tasks before Cloud Manager can be used, including</p> <ul style="list-style-type: none"> ♦ configuring the Cloud Manager Application Server authentication source connections ♦ configuring the Cloud Manager Application Server database connection <p>See Chapter 13, “Using the Cloud Manager Application Server Configuration Tool,” on page 131.</p>

After you have completed the installation and configuration of both the Orchestration components and the Cloud Manager system, continue with [“Setting Up the Cloud Environment”](#) in the [NetIQ Cloud Manager Procedures Guide](#) to start populating your Cloud Manager Application Server and Application Console with components to enable users to provision their own business services.

2 Cloud Manager System Requirements

Before you begin installing the NetIQ Cloud Manager, you need to compare your system resources with the requirements of the product. This section includes information to help you with that evaluation so that you can adequately plan for the installation. The following subsections are included:

- ♦ [Section 2.1, “Cloud Manager Licensing Requirements,” on page 19](#)
- ♦ [Section 2.2, “Cloud Manager Orchestration Server Requirements,” on page 20](#)
- ♦ [Section 2.3, “Cloud Manager Orchestration Agent Requirements,” on page 23](#)
- ♦ [Section 2.4, “Cloud Manager Application Server Requirements,” on page 23](#)
- ♦ [Section 2.5, “Cloud Manager Application Console Requirements,” on page 24](#)
- ♦ [Section 2.6, “Requirements and Cloud Manager Support for the Virtual Environment,” on page 25](#)

2.1 Cloud Manager Licensing Requirements

NetIQ 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.

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 trial key (evaluation) license or a license you purchase from NetIQ. The installation configuration cannot proceed without the license.

2.1.2 Product Evaluation Terms and Procurement

You can download and evaluate NetIQ Cloud Manager without a purchased product license for 90 days. After 90 days, you must purchase a license or discontinue use of the product.

To initiate the evaluation:

- 1 Contact an authorized NetIQ Sales representative at 800-529-3400, or submit a product evaluation request at the [How to Buy NetIQ Cloud Manager \(https://wwwtest.netiq.com/products/cloud-manager/how-to-buy/\)](https://wwwtest.netiq.com/products/cloud-manager/how-to-buy/) Web page.
Your representative will send you a link to the Cloud Manager product download page at the Novell Customer Center.
- 2 On the product download page, click *Get Trial Key* to link to the Product Evaluation Activation page.

- 3 On the Product Evaluation Activation page, select *My Products > Products* to open the Products page.
- 4 On the Products page, select *Novell Cloud Manager* to expand that product list.



- 5 In the *Cloud Manager 2.1* product line, click *Go To* to open the Product Subscription Information page in the Novell Customer Center.
- 6 In the *Downloads* section of the page, select *License* to open the Download page for the trial key.
- 7 Download the trial key file (`trial_keys_3.x_7003_8002-exp-05-01-2013_(xx).txt`) to a location that you can access during the Cloud Manager installation.

After you have obtained the trial key, you can download the product from the Novell Customer Center link sent to you.

For installation and configuration information, see the [NetIQ Cloud Manager 2.x documentation](http://www.novell.com/documentation/cloudmanager2/) (<http://www.novell.com/documentation/cloudmanager2/>) Web site.

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

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

- 1 Stop the Orchestration Server.
`/etc/init.d/netiq-cmosserver stop`
- 2 Copy the purchased license 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

The network machine where you install Cloud Manager Server software must meet the following requirements:

Table 2-1 Orchestration Server Requirements

Item	Requirement
Operating System	One of the following platforms can be used: <ul style="list-style-type: none"> ♦ SUSE Linux Enterprise Server 11 Service Pack 2 (SLES 11 SP2) on the 64-bit (x86-64) architecture (Intel and AMD Opteron processors)
Hardware	<ul style="list-style-type: none"> ♦ Processor: 2.5 GHz 64-bit, or equivalent AMD or Intel processor (minimum); Dual-Core, 2.5 GHz (or greater) 64-bit (recommended) ♦ RAM: 3 GB minimum; 4 GB recommended ♦ Disk Space: 350 MB minimum for installing; 1 GB recommended for managing fewer than 100 resources.

Item	Requirement
Hostname Resolution	The server must resolve device hostnames by using a method such as DNS (recommended).
IP Address	The server must have a static IP address or a permanently leased DHCP address.

Other important requirements you might need to know about the Orchestration Server are included in the following sections:

- ♦ [Section 2.2.1, “Required Network Resources for the Cloud Manager Orchestration Server,” on page 21](#)
- ♦ [Section 2.2.2, “Required Network Resources for the Cloud Manager Orchestration Web Server,” on page 22](#)

2.2.1 Required Network Resources for the Cloud Manager Orchestration Server

The Orchestration Server must allow traffic on TCP ports 80, 8001, 8100, 8101 (these four ports are configurable), and UDP and TCP port 1099 (mandatory).

- ♦ Port 8001 is used for communication with the Administrator Information page.
- ♦ Port 8100 is 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 8101 is also used for invoking the zos command line interface or opening the Java Developer’s toolkit by using TLS.
- ♦ Port 1099 is used with RMI for invoking the zosadmin command line interface or for running the Orchestration Console.

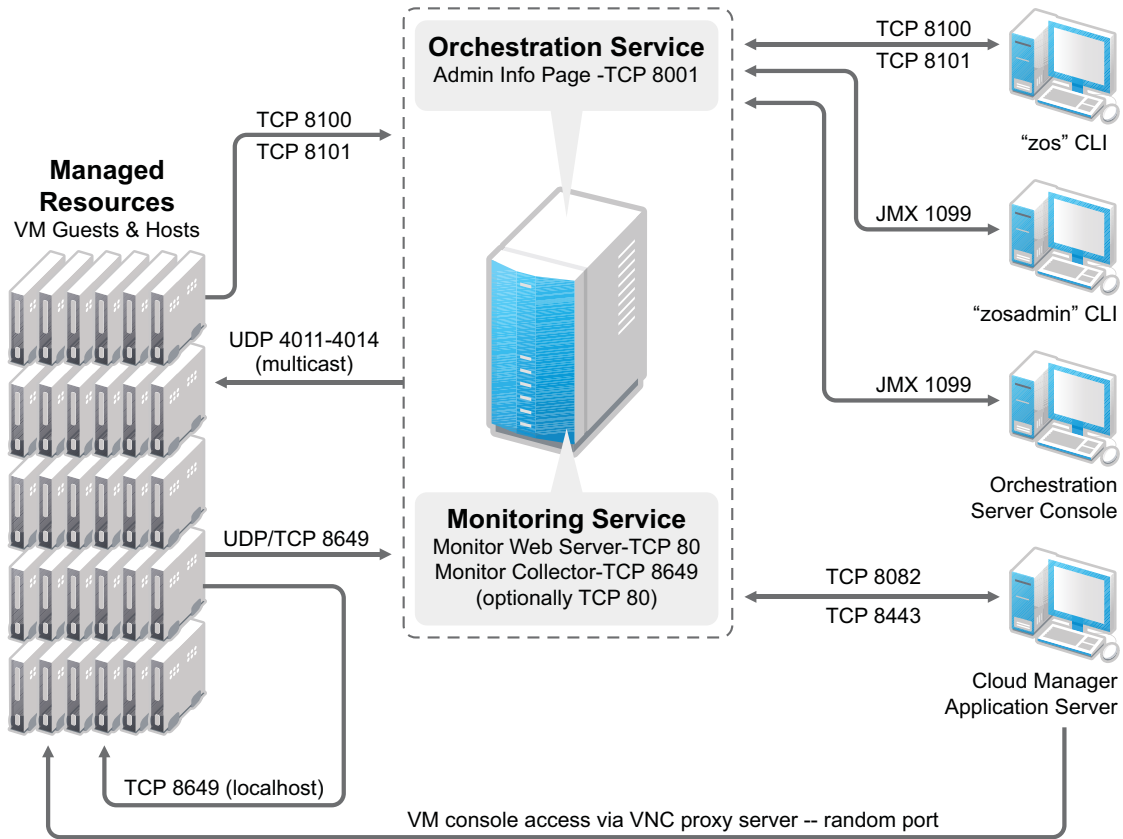
Monitored systems (physical and virtual) send metrics to the Monitoring Server on UDP port 8649. The Monitoring Server is installed on the same system as the Cloud Manager Orchestration Server.

Connections to VM consoles are accommodated through a VNC client. Typically, this means TCP port 5900 for the first VM on a VM host, 5901 for the second, and so on. These connections go to the VM host, exposing the console on behalf of the VM.

Datagrid multicast file transfers use UDP ports 4011-4014. UDP port 4000 is used as a datagrid multicast request port and a control channel port. Multicast groups for datagrid multicast-based file transfers are 239.192.10.10-14.

The following illustration shows these relationships:

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



2.2.2 Required Network Resources for the Cloud Manager Orchestration Web 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 ports 8082 and 8443.

2.3 Cloud Manager Orchestration Agent Requirements

The physical or virtual machine where you install the Orchestration Agent must meet the following minimum requirements:

Table 2-2 Orchestration Agent Requirements

Item	Requirement
Operating System	<p>Linux machines:</p> <ul style="list-style-type: none">♦ SUSE Linux Enterprise Server 10 SP4 (64-bit)♦ SUSE Linux Enterprise Server 11 SP2 (64-bit)♦ Red Hat Enterprise Linux 5 (latest update, 64-bit)♦ Red Hat Enterprise Linux 6 (latest update, 64-bit) <p>Windows machines:</p> <ul style="list-style-type: none">♦ Windows Server 2003 (latest SP, 64-bit)♦ Windows Server 2003 R2 (latest SP, 64-bit)♦ Windows Server 2008 R2 SP1 (64-bit)♦ Windows Server 2008 R2 (latest SP, 64-bit)♦ Windows Server 2008 R2 (latest SP, with HyperV role, 64-bit)
Hardware	The Orchestration Agent does not require a minimum hardware configuration other than a minimum recommended disk space of 100 MB.
TCP Ports	The computing node communicates with the Orchestration Server over a custom protocol. The server listens for the agent on port 8100 and 8101 (port 8101 is for secure agent connections). Network firewalls need to allow outgoing agent connections to these ports on the server.

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.

2.4 Cloud Manager Application Server Requirements

The Cloud Manager Application Server requires the following:

Item	Requirement
Operating System	<p>Any of the following:</p> <ul style="list-style-type: none">♦ SLES 11 SP2 (64-bit): SUSE Linux Enterprise Server 11 Service Pack 2 on the 64-bit (x86-64) architecture (Intel and AMD Opteron processors)

Item	Requirement
Hardware	<p>If the Cloud Manager Application Server is the only application, the following are minimum requirements:</p> <ul style="list-style-type: none"> ♦ Xeon dual-core or higher ♦ 20 GB disk space ♦ 4 GB RAM <p>If the Cloud Manager Application Server and Cloud Manager Orchestration Server are on the same server, the following are minimum requirements:</p> <ul style="list-style-type: none"> ♦ 4 Pentium-class CPU cores ♦ 40 GB disk space ♦ 4 GB RAM
Database	PostgreSQL (included with SLES)
TCP Ports	<p>The following ports are used by the Cloud Manager Application Server. The ports (or their substitutes if not using the defaults) must be open for both inbound and outbound communication:</p> <ul style="list-style-type: none"> ♦ 8061 - ESB HTTP port ♦ 8102 - Karaf SSH port ♦ 8181 - Karaf Management Console port ♦ 8182 - Jetty HTTP default port ♦ 8183 - Jetty HTTPS default port ♦ 10990 - RMI Registry port ♦ 61613 - Active MQ Stomp port ♦ 61616 - Active MQ Openwire port
VNC Ports	<p>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.</p> <p>For more information, see “Configuring Remote Console Access to Workloads” in the <i>NetIQ Cloud Manager Procedures Guide</i>.</p>
LDAP Directory Service	<p>The Cloud Manager Application Server authenticates users by using an LDAP directory. The directory must be either of the following:</p> <ul style="list-style-type: none"> ♦ Microsoft Active Directory ♦ Novell eDirectory

2.5 Cloud Manager Application Console Requirements

The Cloud Manager Application Console is a Web-based application that requires the following:

Item	Requirement
Web Browser	Any of the following: <ul style="list-style-type: none"> ♦ Internet Explorer 9.0 and later: Supported on Windows 7 (64-bit) ♦ Mozilla Firefox 7.x and later: Supported on Windows 7 (64-bit) ♦ Safari 5 and later: Supported on Windows 7 (PC,64-bit)
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

2.6 Requirements and Cloud Manager Support for the Virtual Environment

The following table lists the virtual machine technologies or hypervisors, the host operating system for these technologies, the guest operating systems (also known as virtual machines (VMs) or “workloads”) supported by these technologies, and the provisioning adapter job available in the Cloud Manager Orchestration Server that is used to provision and manage the life cycle of the VMs.

More information about [RHEL 6 VM support](#) in Cloud Manager is also provided in this section.

For more detail about the life cycle management capabilities of Cloud Manager Orchestration, see [Chapter 11, “Configuring Orchestration Provisioning Adapters,” on page 81](#).

Table 2-3 VM Technologies with Supported Host Operating Systems, Guest Operating System, and Provisioning Adapter

Hypervisor or Virtualization Technology	Host Operating System (that is, “VM Hosts”)	Guest Operating System (that is, “VMs” or “Workloads”)	Orchestration Provisioning Adapter
<ul style="list-style-type: none"> ♦ VMware vSphere 4 ♦ VMware vSphere 5 ESXi Only 	Subject to the VMware support matrix	<ul style="list-style-type: none"> ♦ SLES 10 SP3 ♦ SLES 10 SP4 ♦ SLES 11 (latest SP) ♦ RHEL 5 (latest SP) ♦ RHEL 6¹ ♦ Windows Server 2003 R2 (latest SP) ♦ Windows Server 2008 R2 (latest SP) 	vsphere
Citrix XenServer 5.6, latest SP	Citrix XenServer	<ul style="list-style-type: none"> ♦ Windows Server 2008 R2 (latest SP) ♦ Windows Server 2003 R2 (latest SP) ♦ SLES 11 (latest SP) ♦ RHEL 5 (latest SP) ♦ RHEL 6 (latest SP) 	xenserv

Hypervisor or Virtualization Technology	Host Operating System (that is, “VM Hosts”)	Guest Operating System (that is, “VMs” or “Workloads”)	Orchestration Provisioning Adapter
Citrix XenServer 6 Free Edition	Citrix XenServer	<ul style="list-style-type: none"> Windows Server 2008 R2 (latest SP) Windows Server 2003 R2 (latest SP) SLES 11 (latest SP) RHEL 5 (latest SP) RHEL 6 (latest SP) 	xenserv
Microsoft Hyper-V ⁴	Windows Server 2008 R2 with Hyper-V enabled	<ul style="list-style-type: none"> Windows Server 2008 R2 (latest SP) Windows Server 2003 R2 (latest SP) 	hyperv
<ul style="list-style-type: none"> SUSE Xen 4.0 	<ul style="list-style-type: none"> SLES 11 (latest SP) 	<ul style="list-style-type: none"> SLES 10 (latest SP) SLES 11 (latest SP) RHEL 5 (latest SP) RHEL 6 (latest SP)¹ Windows Server 2003 R2 (latest SP)² Windows Server 2008 R2 (latest SP)² 	xen
Kernel-based Virtual Machine for Linux (KVM)	SLES 11 SP1 or SP2 running libvirt 0.7.6 or greater	Subject to the published KVM support matrix (http://www.linux-kvm.org/page/Guest_Support_Status)	kvm

¹ For more information about RHEL 6 VM support, see *RHEL 6 VM Support*, below.

² Windows VMs running on the Xen hypervisor require a VM host CPU with the Intel VT or AMD-V technology available and enabled.

⁴ A complete listing of guest OS support for the Hyper-V hypervisor is available at the [Microsoft TechNet Web site \(http://technet.microsoft.com/en-us/library/cc794868\(WS.10\).aspx\)](http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx) and at the [Windows Server 2008 Hyper-V product page \(http://www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx\)](http://www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx). This matrix shows only those guest OS's supported by Cloud Manager.

RHEL 6 VM Support

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

- Although RHEL uses LVM partitioning by default, we recommend that you do not use it. You need to change the partitioning method manually.
- SLES 11 hosts can mount the ext4 file system if you load the proper kernel module on the host. You can do this by entering the following command at the command line of the SLES 11 host:

```
modprobe -allow-unsupported ext4
```

To allow the ext4 module to be loaded at boot time:

1. Edit the `/etc/modprobe.d/unsupported-modules` file and set `allow_unsupported_modules` to 1.
2. Edit `/etc/sysconfig/kernel` and add ext4 to the `MODULES_LOADED_ON_BOOT` variable.

These procedures work only on SLES 11 kernel, not the SLES 10 kernel.

Making these changes could make the system unavailable for support. The `unsupported-modules` text file states:

“Every kernel module has a ‘supported’ flag. If this flag is not set, loading this module taints your kernel. You will not get much help with a kernel problem if your kernel is marked as tainted. In this case you firstly have to avoid loading of unsupported modules.”

- ♦ Discovered RHEL 6 VMs show appropriate fact values. For example, the value for the `resource.os.type` fact is `rhel6`. The value for `resource.os.vendor.string` is `Red Hat Enterprise Linux Server release 6.0 (Santiago)` and the value for `resource.os.vendor.version` is `6`.
 - ♦ 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 (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.

2.6.1 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 2-4 *Minimum and Recommended Hardware Requirements for VM Host Machines*

Host Operating System	Minimum Requirements	Recommended Hardware
SLES 11 SP2	♦ x86_64	♦ x86_64
	♦ 2 GB RAM	♦ 4+ GB RAM
	♦ 30 GB hard drive space	♦ 100+ GB hard drive space
SLES 11 SP1	♦ x86 or x86_64	♦ x86 or x86_64
	♦ 2 GB RAM	♦ 4+ GB RAM
	♦ 30 GB hard drive space	♦ 100+ GB hard drive space

Host Operating System	Minimum Requirements	Recommended Hardware
Windows Server 2008 R2 enabled with Hyper-V	<ul style="list-style-type: none"> ♦ 1GHz (x86 processor) or 1.4GHz (x64 processor) ♦ 512MB RAM ♦ 10 GB hard drive space 	<ul style="list-style-type: none"> ♦ 2+ GHz ♦ 2+ GB RAM ♦ 40+ GB hard drive space

2.6.2 Supported VMs

The following table lists the virtual machine technologies or hypervisors, the host operating system for these technologies, the guest operating systems (also known as virtual machines (VMs) or “workloads”) supported by these technologies, and the provisioning adapter job available in the Cloud Manager Orchestration Server that is used to provision and manage the life cycle of the VMs.

For more detail about the life cycle management capabilities of Cloud Manager Orchestration, see [Chapter 11, “Configuring Orchestration Provisioning Adapters,” on page 81](#).

Table 2-5 VM Technologies with Supported Host Operating Systems, Guest Operating System, and Provisioning Adapter

Hypervisor or Virtualization Technology	Host Operating System (that is, “VM Hosts”)	Guest Operating System (that is, “VMs” or “Workloads”)	Orchestration Provisioning Adapter
<ul style="list-style-type: none"> ♦ VMware vSphere 4 ♦ VMware vSphere 5 ESXi Only 	Subject to the VMware support matrix	<ul style="list-style-type: none"> ♦ SLES 10 SP3 ♦ SLES 10 SP4 ♦ SLES 11 (latest SP) ♦ RHEL 5 (latest SP) ♦ RHEL 6¹ ♦ Windows Server 2003 R2 (latest SP) ♦ Windows Server 2008 R2 (latest SP) 	vsphere
Citrix XenServer 5.6, latest SP	Citrix XenServer	<ul style="list-style-type: none"> ♦ Windows Server 2008 R2 (latest SP) ♦ Windows Server 2003 R2 (latest SP) ♦ SLES 11 (latest SP) ♦ RHEL 5 (latest SP) ♦ RHEL 6 (latest SP) 	xenserv
Microsoft Hyper-V ⁴	Windows Server 2008 R2 with Hyper-V enabled	<ul style="list-style-type: none"> ♦ Windows Server 2008 R2 (latest SP) ♦ Windows Server 2003 R2 (latest SP) 	hyperv

Hypervisor or Virtualization Technology	Host Operating System (that is, “VM Hosts”)	Guest Operating System (that is, “VMs” or “Workloads”)	Orchestration Provisioning Adapter
♦ SUSE Xen 4.0	♦ SLES 11 (latest SP)	♦ SLES 10 (latest SP) ♦ SLES 11 (latest SP) ♦ RHEL 5 (latest SP) ♦ Windows Server 2003 R2 (latest SP) ² ♦ Windows Server 2008 R2 (latest SP) ²	xen
Kernel-based Virtual Machine for Linux (KVM)	SLES 11 SP1 or SP2 running libvirt 0.7.6 or greater	Subject to the published KVM support matrix (http://www.linux-kvm.org/page/Guest_Support_Status)	kvm

¹ For more information about *RHEL 6 VM Support*, see RHEL 6 VM Support, above.

² S Windows VMs running on the Xen hypervisor require a VM host CPU with the Intel VT or AMD-V technology available and enabled.

⁴ A complete listing of guest OS support for the Hyper-V hypervisor is available at the [Microsoft TechNet Web site \(http://technet.microsoft.com/en-us/library/cc794868\(WS.10\).aspx\)](http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx) and at the [Windows Server 2008 Hyper-V product page \(http://www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx\)](http://www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx). This matrix shows only those guest OS's supported by Cloud Manager.

RHEL 6 VM Support Limitations

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

- ♦ The 64-bit version of RHEL 6, unlike previous versions, does not support installation of the 32-bit `zos-agent*.rpm` package. The install now includes a `novell-zenworks-zos-agent-<version>-${release}.x86_64.rpm` package that should be installed instead. This package is referenced on the `http://server:8001` index page and is included in all agent directories (that is (/RHEL4 /RHEL5 /RHEL6) on the 64-bit distribution CD.
- ♦ SLES 11 hosts can mount the ext4 file system if you load the proper kernel module on the host. You can do this by entering the following command at the command line of the SLES 11 host:

```
modprobe -allow-unsupported ext4
```

To allow the ext4 module to be loaded at boot time:

1. Edit the `/etc/modprobe.d/unsupported-modules` file and set `allow_unsupported_modules` to 1.
2. Edit `/etc/sysconfig/kernel` and add ext4 to the `MODULES_LOADED_ON_BOOT` variable.

These procedures work only on SLES 11 kernel, not the SLES 10 kernel.

Making these changes could make the system unavailable for support. The `unsupported-modules` text file states:

“Every kernel module has a ‘supported’ flag. If this flag is not set, loading this module taints your kernel. You will not get much help with a kernel problem if your kernel is marked as tainted. In this case you firstly have to avoid loading of unsupported modules.”

- ♦ Discovered RHEL 6 VMs show appropriate fact values. For example, the value for the `resource.os.type` fact is `rhel6`. The value for `resource.os.vendor.string` is `Red Hat Enterprise Linux Server release 6.0 (Santiago)` and the value for `resource.os.vendor.version` is `6`. The VM Client has also been modified to show *RHEL 6* as an available OS.
 - ♦ RHEL 6 uses the `udev` service, which testing has shown renames the network interfaces on a cloned VM and causes configuration errors. To turn of 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 (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.

3 Choosing the Installation Packages and Where to Install Them

NetIQ Cloud Manager is comprised of a number of different RPMs that are bundled in different installation patterns, all of which are available on the installation media you [download from Novell \(http://download.novell.com\)](http://download.novell.com), an Attachmate Group associate of NetIQ. 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 RPMs in the install patterns must be installed to a [supported version](#) of SUSE Linux Enterprise Server (SLES) 11. The installation uses the *Add-On Products* utility that is available in SUSE's YaST program.

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

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, "NetIQ Cloud Manager Installation Pattern," on page 31](#)
- [Section 3.2, "Cloud Manager Orchestration Server Install Pattern," on page 32](#)
- [Section 3.3, "Cloud Manager Monitoring Server Pattern," on page 33](#)
- [Section 3.4, "Cloud Manager Orchestration Agent Pattern," on page 34](#)
- [Section 3.5, "Orchestration Console Install Pattern," on page 35](#)
- [Section 3.6, "Monitoring Agent Install Pattern," on page 36](#)

3.1 NetIQ Cloud Manager Installation Pattern

Description: The NetIQ Cloud Manager installation pattern consists of packages for the Cloud Manager Application Server and its Web console. This server communicates with Cloud Manager Orchestration Servers to provide instructions for provisioning, managing, and removing workloads. It also performs user authentication with the LDAP server or Novell Cloud Security Services.

The server requires initial configuration after installation to establish authentication with LDAP, NetIQ Cloud Security Services, or NetIQ Access Manager. The configuration also establishes communication with the Cloud Manager Orchestration Server and its console.

Packages in the Pattern: The table below lists the RPMs in the NetIQ Cloud Manager pattern.

Table 3-1 *NetIQ Cloud Manager Packages*

Install Pattern [Short Name]	Default Packages Installed	Additional Required Patterns	Additional Recommended Patterns	Server /Agent
NetIQ Cloud Manager [cloudmanager]	netiq-cloudmanager postgresql-server			S

When you select the NetIQ Cloud Manager pattern, the `netiq-cloudmanager` and the `postgresql-server` packages are selected by default. Although you would typically install Cloud Manager to use an external PostgreSQL database, selecting the `postgresql-server` package lets you install Cloud Manager to an embedded PostgreSQL server.

You can obtain more information about these patterns and packages in the YaST utility when you have mounted the product ISO.

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 Novell eDirectory.

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

3.2 Cloud Manager Orchestration Server Install Pattern

Description: This 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 server 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: The table below lists the RPMs in the Orchestration Server pattern.

Table 3-2 *Orchestration Server Packages*

Install Pattern [Short Name]	Default Packages Installed	Additional Required Patterns [short name]	Additional Recommended Patterns	Server /Agent
Orchestration Server [zw_zos_server]	netiq-cmos-java	[zw_orch_config]	[zw_mon_server]	S
	novell-pso-ws		[zw_zos_clients]	
	novell-zenworks-orch-config			
	novell-zenworks-orch-config-gui			
	novell-zenworks-zos-server			
	novell-zenworks-zos-server-data-agent			
	novell-zenworks-zos-server-data-clients			
	novell-zenworks-zos-server-data-jre			
	novell-zenworks-zos-server-data-livedcd			

You can obtain more information about these patterns in the YaST utility when you have mounted the product ISO.

NOTE: Orchestration Server patterns are labeled version 3.2.0 in the NetIQ Cloud Manager 2.2.0 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, we strongly recommend 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 Novell 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.

Further, be advised that installing the server on a VM slows down the performance of the product.

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

3.3 Cloud Manager Monitoring Server Pattern

Description: The Cloud Manager Monitoring Server is an Apache Web server that uses open source Ganglia monitors 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 text interface at the Linux console (`./config`) or a GUI configuration wizard (`./guiconfig`).

Packages in the Pattern: The table below lists the RPMs in the Orchestration Server pattern.

Table 3-3 *Orchestration Server Packages*

Install Pattern [Short Name]	Default Packages Installed	Additional Required Patterns [short name]	Additional Recommended Patterns	Server /Agent
Monitoring Server	libconfuse0	[zw_mon_agent]		S
[zw_mon_server]	novell-zenworks-monitor-gmetad	[zw_orch_config]		
	novell-zenworks-monitor-web			
	novell-zenworks-orch-config			
	novell-zenworks-orch-config-gui			

You can obtain more information about these patterns and packages in the YaST utility when you have mounted the product ISO.

NOTE: Monitoring Server patterns are all version 3.2.0 in the NetIQ Cloud Manager 2.2.0 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 Pattern

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`).

You can also install the agent from a Windows installation program or use the Linux pattern files to install to RHEL machines. For more information, see [Section 5.2, “Alternative Installation Methods for the Orchestration Agent,”](#) on page 45.

Packages in the Pattern: The table below lists the RPMs in the Orchestration Agent pattern.

Table 3-4 *Orchestration Server Packages*

Install Pattern [Short Name]	Default Packages Installed	Additional Required Patterns [short name]	Additional Recommended Patterns	Server /Agent
Orchestration Agent [zw_zos_agent]	cabextract chntpw fuse netiq-cmos-java novell-zenworks-orch-config novell-zenworks-orch-config-gui novell-zenworks-zos-agent ntfs-3g	[zw_orch-config]	[zw_mon_agent]	A

You can obtain more information about these patterns inside the YaST utility when you have mounted the product ISO.

NOTE: Orchestration Agent patterns are labeled version 3.2.0 in the NetIQ Cloud Manager 2.2.0 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, NetIQ recommends you deploy an Orchestration Agent on a dedicated system. For more information, see “[The VMware vSphere Provisioning Adapter](#)” in the [NetIQ Cloud Manager Component Reference](#).

3.5 Orchestration Console Install Pattern

Description: The Cloud Manager Orchestration Server Console is a java-based thick client that administers the functionality of the Orchestration Server from any SLES 11 server or a Windows 7 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 [NetIQ Cloud Manager Component Reference](#).

Packages in the Pattern: The table below lists the RPMs in the Orchestration Console pattern.

Table 3-5 *Orchestration Console Packages*

Install Pattern [Short Name]	Default Packages Installed	Additional Required Patterns [short name]	Additional Recommended Patterns	Server /Agent
Orchestration Agent	netiq-cmos-java novell-zenworks-zos-clients [zw_zos_clients]			S or A

You can obtain more information about these patterns inside the YaST utility when you have mounted the product ISO.

NOTE: The Orchestration Console pattern is labeled version 3.1.4 in the NetIQ Cloud Manager 2.12.1.4 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 52.

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.

3.6 Monitoring Agent Install Pattern

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).

You can also install the agent from a Windows installation program or use the Linux pattern files to install to RHEL machines. For more information, see [Section 5.4, “Alternative Installation Methods for the Cloud Manager Monitoring Agent,”](#) on page 55.

Packages in the Pattern: The table below lists the RPMs in the Monitoring Agent pattern.

Table 3-6 *Orchestration Server Packages*

Install Pattern [Short Name]	Default Packages Installed	Additional Required Patterns [short name]	Additional Recommended Patterns	Server /Agent
Monitoring Agent	libconfuse0	[zw_orch_config]		A
[zw_mon_agent]	novell-zenworks-monitor-gmond novell-zenworks-orch-config novell-zenworks-orch-config-gui			

You can obtain more information about these patterns and packages in the YaST utility when you have mounted the product ISO.

NOTE: Monitoring Agent patterns are all version 3.2.0 in the NetIQ Cloud Manager 2.2.0 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 55.

4 Orchestration Components Preinstallation Tasks

Before you install the NetIQ Cloud Manager Orchestration components, you need to prepare the environment where those are to be installed:

- ♦ [Section 4.1, “Gathering Certificate and License Information,” on page 39](#)
- ♦ [Section 4.2, “Preparing the Server When Multiple NICs and DNS Addresses Exist,” on page 39](#)

4.1 Gathering Certificate and License Information

Before you install NetIQ 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 NetIQPartner for purchase information.

- ♦ (Optional) Certificate authority information (internal, or signed certificate, private key, and public certificate).

4.2 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 (127.0.0.1 or 127.0.0.2) address 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 hostname is found to be a loopback address (for example, 127.0.0.2), 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 Component Installation

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

- ♦ [Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 43](#)
- ♦ [Chapter 6, “Installing Cloud Manager Application Server Components,” on page 59](#)

5 Installing Cloud Manager Orchestration Components

The RPMs in the Orchestration install patterns must be installed to a [supported version](#) of SUSE Linux Enterprise Server (SLES) 11.

Some Cloud Manager RPMs 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.

- ♦ [Section 5.1, “SLES 11 Standard Installation,” on page 44](#)
- ♦ [Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 45](#)
- ♦ [Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 52](#)
- ♦ [Section 5.4, “Alternative Installation Methods for the Cloud Manager Monitoring Agent,” on page 55](#)

After the initial installation and configuration, installers for some Cloud Manager Orchestration components for other operating systems become available in the Orchestration file system. For more information about these alternative post-installation methods, see [Section 5.2, “Alternative Installation Methods for the Orchestration Agent,” on page 45](#), [Section 5.3, “Alternative Installation Methods for the Orchestration Console and Clients,” on page 52](#), and [Section 5.4, “Alternative Installation Methods for the Cloud Manager Monitoring Agent,” on page 55](#).

NetIQ recommends that you install and configure the Cloud Manager Orchestration components before continuing with the installation and configuration of Cloud Manager components. For more information, see the [Chapter 6, “Installing Cloud Manager Application Server Components,” on page 59](#).

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 45](#).

For information about uninstalling Cloud Manager components, see [Chapter 22, “Uninstalling Orchestration Component Patterns from a SLES Server,” on page 249](#).


For advanced configuration tasks and methods for optimizing the Cloud Manager Orchestration components, see the [NetIQ Cloud Manager Component Reference](#).

5.1 SLES 11 Standard Installation

The steps for installing Cloud Manager Orchestration components on a SLES 11 server, including the Orchestration Server, Orchestration Agent, the Orchestration Console (accompanied by other Orchestration clients), and the Cloud Manager Monitoring Server and Monitoring Agent are included in this section.

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 Packages and Where to Install Them,”](#) on page 31 can help you make that decision.

- 1 Log in to the target SLES server as `root`, then open YaST or YaST2. You should install the Orchestration Server on a dedicated server for optimal performance.
- 2 Download the appropriate NetIQ Cloud Manager ISO to the SLES server.
or
Load the NetIQ Cloud Manager DVD on the SLES server.
- 3 Define the NetIQ 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 box.



- 5 Select the installation pattern that contains the Orchestration component packages you want to install on this server.
You should have previously decided which packages to install. For more information, see [Chapter 3, “Choosing the Installation Packages and Where to Install Them,”](#) on page 31.
- 6 Click *OK* to install the packages.
- 7 When package installation is complete, click *OK* to close the Installed Add-On Products dialog box.

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 either the Cloud Manager Web console or a Cloud Manager Mobile Client. This section includes information about the installation methods you can use that differ from the standard installation on a SLES machine.

The alternative agent installation methods vary depending on the platform you are installing to. You can install the agent on most SLES 11 servers, on most RHEL 5 or RHEL 6 servers, on most Windows 2003 or 2008 servers, on most Windows (XP, Windows 7, or Vista) desktops, or on the SLED 11 SP1 desktop. For exact requirements, see [Section 2.3, “Cloud Manager Orchestration Agent Requirements,” on page 23](#).

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 49](#)

Windows Installation Source: The Windows installation program for the agent is located on the install media at `\Windows\zosagent_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 53](#).

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 46](#).

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 56](#).

Linux Installation Source: The manual installation procedure for the agent 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 5.3.2, “Installing the Console and Clients on Windows,” on page 53](#).

This section includes the following information:

- [Section 5.2.1, “Obtaining the Agent Installer and Supporting Files from the Administrator Information Page,” on page 46](#)
- [Section 5.2.2, “Installing the Agent on Windows Machines,” on page 47](#)
- [Section 5.2.3, “Manually Installing the Agent Packages on SLES Machines,” on page 47](#)
- [Section 5.2.4, “Manually Installing the Agent Linux Packages on RHEL Machines,” on page 48](#)
- [Section 5.2.5, “Advanced Agent Installation Methods,” on page 49](#)

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

- ♦ Internet Explorer, version 6.0 or higher
- ♦ Netscape Navigator, version 6.0 or higher
- ♦ Firefox, version 1.5 or higher

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
<u>Microsoft Windows Server</u> <ul style="list-style-type: none">▪ Windows Server 2003 (latest SP, 64-bit)▪ Windows Server 2003 R2 (latest SP, 64-bit)▪ Windows Server 2008 R2 SP1 (64-bit)▪ Windows Server 2008 R2 Hyper-V (latest SP, 64-bit)	zosagent_windows_3.2.0_with_jre.exe	
<u>SUSE Linux Enterprise Server (SLES) RPM</u> <ul style="list-style-type: none">▪ SLES 10 SP4 (64-bit)▪ SLES 11 SP2 (64-bit)		novell-zenworks-zos-agent-3.2.0-232580.x86_64.rpm (also requires Cloud Manager Orchestration Server Java RPM) Java 1.7.0 (64-bit)
<u>Red Hat Enterprise Linux Server (RHEL) RPM</u> <ul style="list-style-type: none">▪ RHEL 5 (latest update, 64-bit)▪ RHEL 6 (latest update, 64-bit)		novell-zenworks-zos-agent-3.2.0-232580.x86_64.rpm (also requires Cloud Manager Orchestration Server Java RPM) Java 1.7.0 (64-bit)
<u>Community Enterprise Operating System (CentOS) RPM</u> <ul style="list-style-type: none">▪ CentOS 5.8 (latest update, 64-bit)▪ CentOS 6.0 (latest update, 64-bit)		novell-zenworks-zos-agent-3.2.0-232580.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
<u>Microsoft Windows</u> <ul style="list-style-type: none">▪ Windows 7 (latest SP, 64-bit)	zosclients_windows_3.2.0_with_jre.exe	
<u>SUSE Linux Enterprise Server (SLES)</u> <ul style="list-style-type: none">▪ SLES 11 SP2 (64-bit)		novell-zenworks-zos-clients-3.2.0-232580.i386.rpm (also requires Cloud Manager Orchestration Server Java RPM)

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.

Use the following steps to install the agent on a Windows computing resource:

- 1 At the location where you copied the Windows agent installer file (`zosagent_windows_3_2_0_with_jre.exe`), double-click the filename to run the installer.
When you launch the installer on Windows XP or 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.
- 2 Click *Next* to display the Select Destination Directory dialog box.
- 3 Accept the default location, then click *Next* to display the Select Start Menu Folder page of the Setup Wizard.
- 4 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.
- 5 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.
- 6 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, make sure you select the server you previously associated with this agent.

- 7 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.
- 8 Click *Next* to run the Orchestration Agent installation until the Agent Setup Wizard completion page is displayed:.
- 9 Click *Finish* to exit the setup.
- 10 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-3.2.0-<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-3.2.0-<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 machines, the information in this section can help you manually install those files on RHEL 5 or RHEL 6.

- [“Required Agent Installation Files for RHEL Machines” on page 48](#)
- [“Manually Installing Orchestration Agents on RHEL 5” on page 49](#)
- [“Manually Installing Orchestration Agents on RHEL 6” on page 49](#)

Required Agent Installation Files for RHEL Machines

The table below lists Orchestration Agent packages that you need to install on RHEL 5 or RHEL 6 servers. You can find them on the downloaded 32-bit or 64-bit ISO in the `/RHEL5` or `/RHEL6` directories.

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

Platform	Installation Package Name
RHEL 5 (32-bit)	<code>novell-zenworks-orch-config-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-orch-config-gui-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-zos-agent-3.2.0-<build_number>.i586.rpm</code>
	<code>novell-zenworks-zos-java-1.6.0_sun_update14-1.i586.rpm</code>
RHEL 5 (64-bit)	<code>novell-zenworks-orch-config-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-orch-config-gui-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-zos-agent-3.2.0-<build_number>.x86_64.rpm</code>
	<code>netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm</code>
RHEL 6 (32-bit)	<code>novell-zenworks-orch-config-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-orch-config-gui-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-zos-agent-3.2.0-<build_number>.i586.rpm</code>
	<code>novell-zenworks-zos-java-1.6.0_sun_update14-1.i586.rpm</code>
RHEL 6 (64-bit)	<code>novell-zenworks-orch-config-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-orch-config-gui-3.2.0-<build_number>.noarch.rpm</code>
	<code>novell-zenworks-zos-agent-3.2.0-<build_number>.x86_64.rpm</code>
	<code>netiq-cmos-java-1.7.0_sun_update9-3.x86_64.rpm</code>

Manually Installing Orchestration Agents on RHEL 5

To install the four packages required for the Orchestration Agent on RHEL 5:

- 1 Download the pertinent 32-bit or 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 11 ISO, the command is:

```
$ mount -o loop NetIQ_Cloud_Manager-3.2.0-SLE11.x86_64.iso /mnt
```
- 3 Change your working directory to the location of the RHEL package:

```
$ cd /mnt/RHEL5
```
- 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 66](#) for an explanation of the configuration for the Orchestration Agent.

Manually Installing Orchestration Agents on RHEL 6

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

- 1 Download the pertinent 32-bit or 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 11 ISO, the command is:

```
$ mount -o loop NetIQ_Cloud_Manager-3.2.0-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 66](#) 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 50](#)
- ♦ [“Using an Orchestration Job to Install the Orchestration Agent on a VM Host” on page 51](#)
- ♦ [“Automatically Installing the Agent on a VM” on page 52](#)

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 50
- ◆ “Silent Installation and Configuration of the Orchestration Agent RPM” on page 51

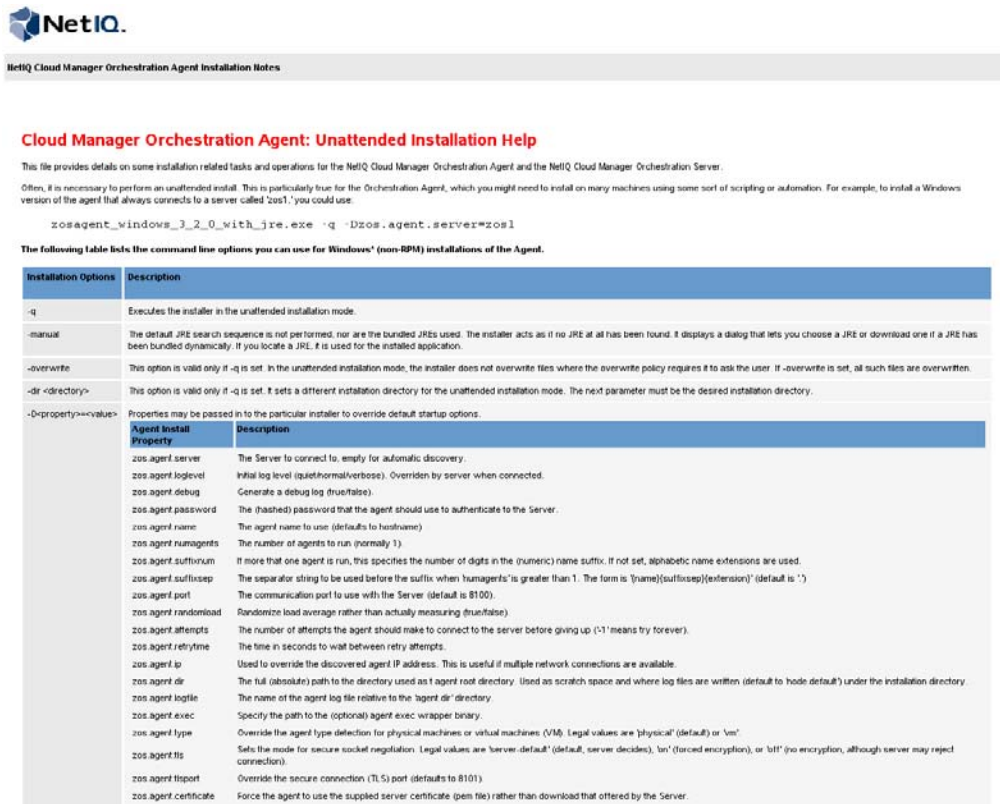
Silent Install and Configuration of the Orchestration Agent for Windows

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



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_2_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.																																										
-Dproperty=<value>	Properties may be passed in to the particular installer to override default startup options.																																										
<table border="1"><thead><tr><th>Agent Install Property</th><th>Description</th></tr></thead><tbody><tr><td>zos.agent.server</td><td>The Server to connect to, empty for automatic discovery.</td></tr><tr><td>zos.agent.loglevel</td><td>Initial log level (quiet/normal/verbose). Overridden by server when connected.</td></tr><tr><td>zos.agent.debug</td><td>Generate a debug log (true/false).</td></tr><tr><td>zos.agent.password</td><td>The (hashed) password that the agent should use to authenticate to the Server.</td></tr><tr><td>zos.agent.name</td><td>The agent name to use (defaults to hostname).</td></tr><tr><td>zos.agent.numagents</td><td>The number of agents to run (normally 1).</td></tr><tr><td>zos.agent.suffixnum</td><td>If more than one agent is run, this specifies the number of digits in the (numeric) name suffix. If not set, alphabetic name extensions are used.</td></tr><tr><td>zos.agent.suffixsep</td><td>The separator string to be used before the suffix when 'numagents' is greater than 1. The form is {name}{suffixsep}{extension} (default is ':').</td></tr><tr><td>zos.agent.port</td><td>The communication port to use with the Server (default is 8100).</td></tr><tr><td>zos.agent.randomload</td><td>Randomize load average rather than actually measuring (true/false).</td></tr><tr><td>zos.agent.attempts</td><td>The number of attempts the agent should make to connect to the server before giving up (-1 means try forever).</td></tr><tr><td>zos.agent.retrytime</td><td>The time in seconds to wait between retry attempts.</td></tr><tr><td>zos.agent.ip</td><td>Used to override the discovered agent IP address. This is useful if multiple network connections are available.</td></tr><tr><td>zos.agent.dir</td><td>The full (absolute) path to the directory used as a temporary directory. Used as scratch space and where log files are written (default to %home%\default) under the installation directory.</td></tr><tr><td>zos.agent.logfile</td><td>The name of the agent log file relative to the 'agent.dir' directory.</td></tr><tr><td>zos.agent.exec</td><td>Specify the path to the (optional) agent exec wrapper binary.</td></tr><tr><td>zos.agent.type</td><td>Override the agent type detection for physical machines or virtual machines (VM). Legal values are 'physical' (default) or 'vm'.</td></tr><tr><td>zos.agent.tls</td><td>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).</td></tr><tr><td>zos.agent.tlsport</td><td>Override the secure connection (TLS) port (defaults to 8101).</td></tr><tr><td>zos.agent.certificate</td><td>Force the agent to use the supplied server certificate (pem file) rather than download that offered by the Server.</td></tr></tbody></table>		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). Overridden 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 than 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 a temporary directory. Used as scratch space and where log files are written (default to %home%\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.
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). Overridden 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 than 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 a temporary directory. Used as scratch space and where log files are written (default to %home%\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.																																										

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 (referenceto Agent Install).
- 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 RPMs as described in [Section 5.2.3, “Manually Installing the Agent Packages on SLES Machines,” on page 47](#). Make sure 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 any SLES 10 or SLES 11 server, or SLED 11 or a Windows (XP, Windows 7, or Vista) desktop.

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 53](#).

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 54](#).

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:

- ♦ Internet Explorer, version 6.0 or higher
- ♦ Netscape Navigator, version 6.0 or higher
- ♦ Firefox, version 1.5 or higher

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

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 <ul style="list-style-type: none"> ■ Windows Server 2003 (latest SP, 64-bit) ■ Windows Server 2003 R2 (latest SP, 64-bit) ■ Windows Server 2008 R2 SP1 (64-bit) ■ Windows Server 2008 R2 Hyper-V (latest SP, 64-bit) 	zosagent_windows_3_2_0_with_jre.exe	
SUSE Linux Enterprise Server (SLES) RPM <ul style="list-style-type: none"> ■ SLES 10 SP4 (64-bit) ■ SLES 11 SP2 (64-bit) 		novell-zenworks-zos-agent-3.2.0-232580.x86_64.rpm (also requires Cloud Manager Orchestration Server Java RPM) Java 1.7.0 (64-bit)
Red Hat Enterprise Linux Server (RHEL) RPM <ul style="list-style-type: none"> ■ RHEL 5 (latest update, 64-bit) ■ RHEL 6 (latest update, 64-bit) 		novell-zenworks-zos-agent-3.2.0-232580.x86_64.rpm (also requires Cloud Manager Orchestration Server Java RPM) Java 1.7.0 (64-bit)
Community Enterprise Operating System (CentOS) RPM <ul style="list-style-type: none"> ■ CentOS 5.8 (latest update, 64-bit) ■ CentOS 6.0 (latest update, 64-bit) 		novell-zenworks-zos-agent-3.2.0-232580.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 <ul style="list-style-type: none"> ■ Windows 7 (latest SP, 64-bit) 	zosclients_windows_3_2_0_with_jre.exe	
SUSE Linux Enterprise Server (SLES) <ul style="list-style-type: none"> ■ SLES 11 SP2 (64-bit) 		novell-zenworks-zos-clients-3.2.0-232580.i686.rpm (also requires Cloud Manager Orchestration Server Java RPM)

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

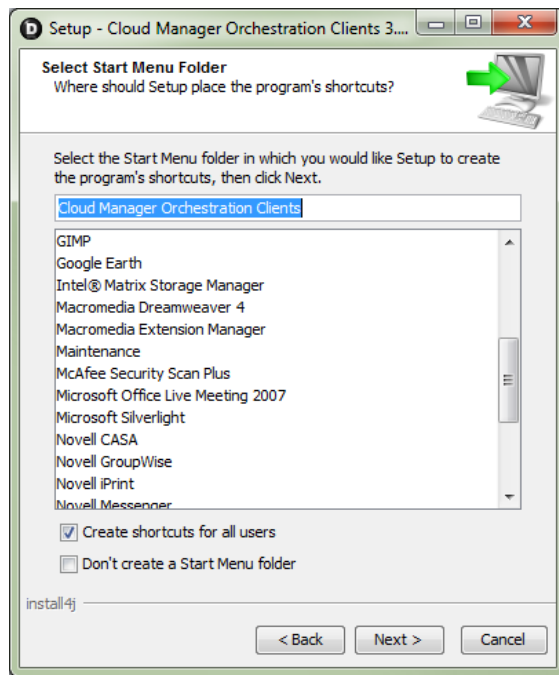
5.3.2 Installing the Console and Clients on Windows

- 1 At a Windows XP, Windows 7, or Windows Vista location where you downloaded the client installer file for Windows, double-click the `zosclients_windows_3_2_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 [NetIQ Cloud Manager Component 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 [NetIQ Cloud Manager Component 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-3.2.0-*build_number*.i586.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-3.2.0-<build_number>.i586.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 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.

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 55](#)
- ♦ [Section 5.4.2, “Installing the Cloud Manager Monitoring Agent On Windows Machines,” on page 56](#)

For more information about agent installation, see [Appendix 5.2.5, “Advanced Agent Installation Methods,” on page 49](#).

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 55](#)

Cloud Manager Monitoring Agent Installation Files for Linux Servers

The Cloud Manager Monitoring Agent uses both the agent program files

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

Operating System	Installation File
♦ SLES 11 SP2 (64-bit)	♦ <cd>/suse/setup/descr/zw_mon_agent-3.1.3-0.x86_64.pat
♦ RHEL 6 (latest update, 64-bit)	♦ <cd>/RHEL6/novell-zenworks-monitor-gmond-3.0.4-67.1.x86_64.rpm
♦ RHEL 5 (latest update, 64-bit)	♦ <cd>/RHEL5/novell-zenworks-monitor-gmond-3.0.4-67.1.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 47](#)) 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 2008 R2 latest SP (with HyperV role, 64-bit)	♦ <cd>/Windows/GmondSetup.exe
♦ Windows 2008 R2 latest SP (64-bit)	♦ <cd>/Windows/GmondSetup.exe
♦ Windows 2003 latest SP (64-bit)	♦ <cd>/Windows/GmondSetup.exe
♦ Windows 2003 R2 latest SP (64-bit)	♦ <cd>/Windows/GmondSetup.exe

This section includes the following information:

- ♦ [“Installing the Cloud Manager Monitoring Agent for Windows” on page 56](#)
- ♦ [“Configuring the Monitoring Agent for Windows” on page 57](#)

Installing the Cloud Manager Monitoring Agent for Windows

- ♦ [“Hardware and Software Requirements” on page 56](#)
- ♦ [“Installing the Monitoring Agent” on page 56](#)
- ♦ [“Starting and Stopping the Monitoring Agent” on page 57](#)
- ♦ [“Uninstalling the Monitoring Agent” on page 57](#)

Hardware and Software Requirements

The Cloud Manager Monitoring Agent (gmond) can be installed only on Windows Server 2003 or Windows Server 2008 machines 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 local Windows Server 2003 or Windows Server 2008 machine.

- 1 Download the pertinent Add-On ISO from the product DVD.
- 2 Create a CD from the ISO or use ISO Buster (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, click *Accept* in the User Account Control dialog box to allow the installation to proceed as system administrator.

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 Novell.

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

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

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

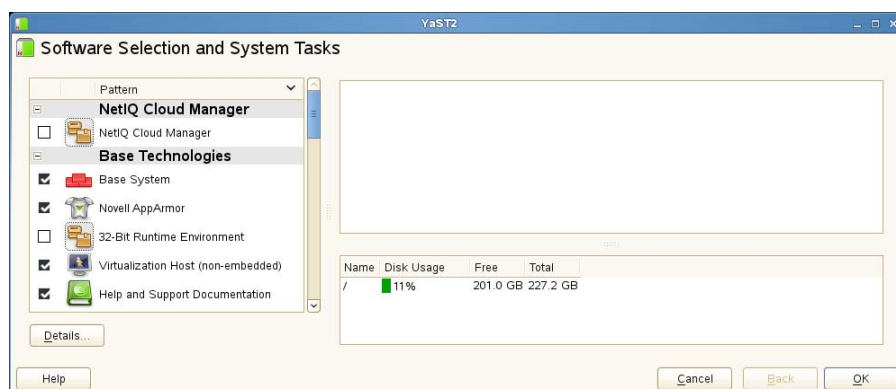
- ♦ [Section 6.1, “Installing to SLES 11,” on page 59](#)

6.1 Installing to SLES 11


The RPMs in the Cloud Manager install patterns must be installed to a [supported version](#) of SUSE Linux Enterprise Server (SLES) 10 or 11. 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 19](#).

Some Cloud Manager RPMs 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 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.

- 1 Log in to the target SLES server as `root`, then open YaST.
- 2 From the NetIQ product downloads Web site, download the appropriate NetIQ Cloud Manager ISO to the SLES server.
or
Load the NetIQ Cloud Manager DVD on the SLES server.
- 3 Define the NetIQ 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 box.



- 5 Select the *NetIQ 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 box.



Standard Cloud Manager Component Configuration

- ♦ [Chapter 7, “Configuring Cloud Manager Orchestration Components,” on page 63](#)
- ♦ [Chapter 8, “Configuring Connections to the Cloud Manager Application Server,” on page 69](#)
- ♦ [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 Orchestration Provisioning Adapters,” on page 81](#)
- ♦ [Chapter 12, “Configuring Sysprep or Autoprep,” on page 111](#)
- ♦ [Chapter 13, “Using the Cloud Manager Application Server Configuration Tool,” on page 131](#)

7 Configuring Cloud Manager Orchestration Components

This section discusses the basic configuration of all NetIQ 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, make sure 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 3.0. This might cause confusion because they are to be used with Cloud Manager Application components, which are version 2.0.

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 64](#)
- ♦ [Section 7.2, “Configuring the Monitoring Server and Monitoring Agent,” on page 66](#)
- ♦ [Section 7.3, “Configuring the Orchestration Agent,” on page 66](#)
- ♦ [Section 7.4, “Validating and Optimizing the Orchestration Installation,” on page 68](#)

When the installation and configuration are complete, you need to [validate and optimize](#) the configuration. You can then 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 Make sure you are ready with the information that you’ll be prompted for during the configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Configuration Type	<p>Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation.</p> <p>This section discusses standard installation, so specify <code>s</code> (for <code>standard</code>) or press Enter to accept the default. For more information about High Availability configuration, see the Chapter 14, “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support,” on page 155.</p>
Server IP Address	<p>You need to know the IP address of this server if you have multiple network interfaces. If specified, the server uses the specified hostname or IP address for binding RMI connections.</p>
Grid Name	<p>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.</p> <p>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.</p>
Administrator User	<p>The name you specify here is required when you access the Orchestration Console or the <code>zosadmin</code> command line interface.</p> <p>You should remember this username for future logins.</p>
Administrator Password	<p>The password you specify here is required when you access the Orchestration Console or the <code>zosadmin</code> command line interface.</p> <p>You should remember this username for future logins.</p>

Server Configuration Requirement	Explanation and Action
Auditing Database JDBC URL	<p>If you answer <i>yes</i> 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.</p> <p>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.</p> <p>For more information, see Chapter 16, "Configuring the Orchestration Server to Use an Audit Database," on page 177.</p>
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 Novell, 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	<p>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.</p> <p>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.</p> <p>This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the <i>Configure Advanced Settings</i> 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.</p>
(Optional) Path to TLS Server Certificate and TLS Server Private Key	<p>A PEM-encoded TLS certificate and key is needed for secure communication between the Orchestration Server and Orchestration Agent.</p> <p>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.</p> <p>This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the <i>Configure Advanced Settings</i> check box in the wizard, this parameter is listed, but default values are provided only if the previous value is manually set to <i>no</i>.</p>

- 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 Make sure 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?	<p>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.</p> <p>The configuration lets you choose not to configure this computer as a Monitoring Server.</p>
Hostname or IP Address of the Monitoring Server	You need to know the hostname 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 Make sure you are ready with the information that you'll be prompted for during the Monitoring Server configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Configuration Type	<p>Your answer here determines whether this iconfiguration takes place on a standard installation or on a High Availability installation.</p> <p>This section discusses standard installation, so specify <code>s</code> (for <code>standard</code>) or press Enter to accept the default. For more information about High Availability configuration, see the Chapter 14, “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support,” on page 155.</p>
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?	<p>The Agent relies on the Orchestration Server's TLS certificate as verification that it is communicating with the correct Orchestration Server.</p> <p>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.</p>
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	<p>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.</p> <p>For an Agent installed on ESX, configure port 8101.</p>
(Optional) Agent IP Address	<p>You can specify a local bind address for the agent if you want to.</p> <p>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.</p>
Path to Server Certificate	<p>Specify the path to the Orchestration Server certificate file. The default path is <code>/root/zos_server_cert.pem</code>.</p> <p>NOTE: This configuration parameter is considered an advanced setting for the Orchestration Agent in the GUI Configuration Wizard, but only if you set <i>Provide Existing Orchestration Server Certificate</i> to <i>yes</i>.</p>

- 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 50](#).

7.4 Validating and Optimizing the Orchestration Installation

- 1 Open the configuration log file (`/var/opt/novell/novell_zenworks_orch_install.log`) to make sure 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 47](#).

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 47](#).

- 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.

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 69](#)
- ♦ [Section 8.2, “Enabling a Non-Secure Connection,” on page 70](#)

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 69](#)

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
```

8.2 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 Server 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 Server 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 box.

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 box, specify the IP address of the Orchestration Server in the *Server Address* field, then click *OK* to display the login dialog box.
- 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 75](#)
- ♦ [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](#)

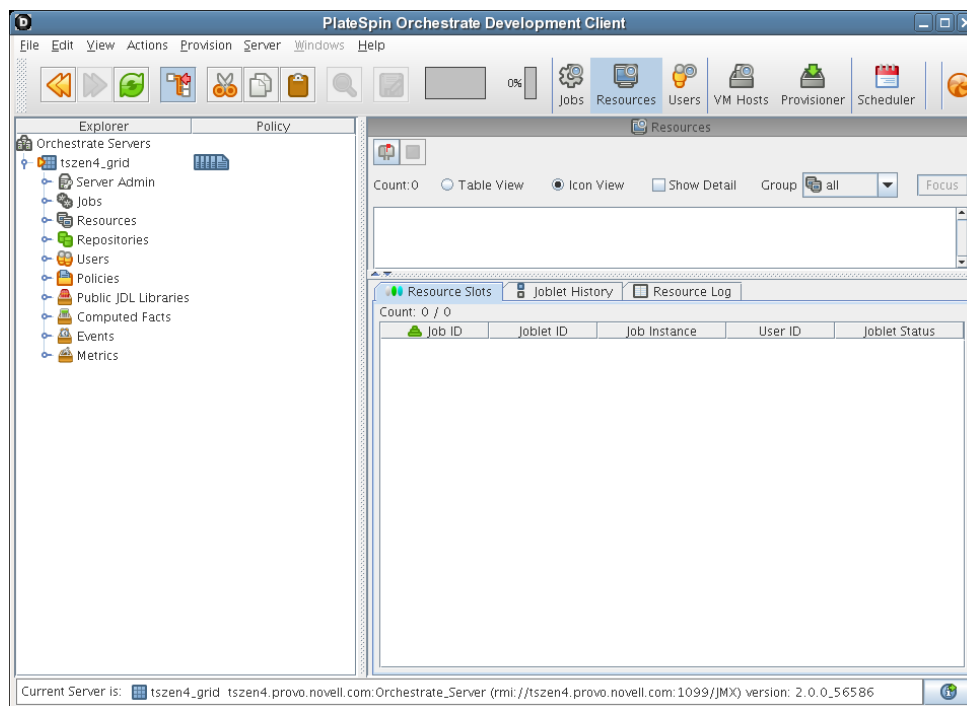
When resources are created, connected to the Orchestration Server and online, a provisioning adapter job deploys and runs a discovery process on its own.

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. For more detailed information about the Resources Monitor, see “Monitoring Server Resources” in the [NetIQ Cloud Manager Component Reference](#).

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 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 box.

This dialog box 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 Make sure that the *Auto Register Agents* check box on the grid object's *Authentication* page is not selected (see [Step 2 on page 75](#)).
- 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 box.
 - 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 box.



This dialog box 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 Make sure *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 46.
- 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 <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> Windows 7 (latest SP; 32-bit and 64-bit) 	zosclients_windows_3_0_0_with_jre.exe	
SUSE Linux Enterprise Server (SLES) RPM <ul style="list-style-type: none"> 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_3_2_0_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 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_3_2_0_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: Make sure 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 Orchestration Provisioning Adapters

You can complete the configuration of the Cloud Manager Orchestration system by preparing it to receive information about the VMs it discovers in your installed hypervisor environment. This VM discovery is made possible when you configure the Orchestration provisioning adapter jobs for each hypervisor environment. This section discusses how to configure the pre-packaged provisioning adapters to discover VMs.

- ♦ [Section 11.1, “Configuring the vSphere Provisioning Adapter,” on page 81](#)
- ♦ [Section 11.2, “Configuring the Citrix XenServer Provisioning Adapter,” on page 95](#)
- ♦ [Section 11.3, “Configuring the Hyper-V Provisioning Adapter,” on page 101](#)
- ♦ [Section 11.4, “Configuring the SUSE Xen Provisioning Adapter,” on page 103](#)
- ♦ [Section 11.5, “Configuring the KVM Provisioning Adapter,” on page 107](#)

11.1 Configuring the vSphere Provisioning Adapter

This section includes the following information:

- ♦ [Section 11.1.1, “Configuring the vSphere Provisioning Adapter to Discover VMs,” on page 81](#)
- ♦ [Section 11.1.2, “Discovering Enterprise Resources in Multiple vSphere Environments,” on page 92](#)

For additional information about the vSphere provisioning adapter, see [“The VMware vSphere Provisioning Adapter”](#) in the *NetIQ Cloud Manager Component Reference*.

11.1.1 Configuring the vSphere Provisioning Adapter to Discover VMs

The content of this section includes information to help you configure the vsphere provisioning adapter job, which authenticates to a VMware hypervisor environment and then discovers VMs in that environment.

The information is organized in the following sections:

- ♦ [“Initial Configuration” on page 82](#)
- ♦ [“Policy Configuration Summary for the vSphere Provisioning Adapter” on page 83](#)
- ♦ [“Assigning a vSphere VM to a Resource Pool” on page 84](#)
- ♦ [“Setting Up Orchestration VNC for a VM Managed by vSphere” on page 85](#)
- ♦ [“Setting Up Orchestration to Accommodate VMware DRS Clustering and Updates” on page 87](#)
- ♦ [“Constraining vSphere VMs to Their Assigned Resource Pools” on page 90](#)
- ♦ [“Configuring Cloud Manager Orchestration for a vSphere MetroCluster Environment” on page 90](#)

Initial Configuration

Before you can provision and manage VMs with the vsphere provisioning adapter job, you must perform some initial steps to configure it in order to get it running.

- 1 Make sure that the Orchestration Agent is installed and started on a supported host.

For more information, see [Chapter 5, “Installing Cloud Manager Orchestration Components,” on page 43](#).

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 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 box.
 - 4b Fill in the fields of the dialog box:
 - ♦ **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 [NetIQ Cloud Manager Component 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 box.
 - 8a Fill in the fields of the dialog box:
 - ♦ **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 Webservice 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 could also 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 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*.
- 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 box, 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 *NetIQ Cloud Manager Component Reference*.
- If you want to connect multiple vCenter Servers, refer to [Section 11.1.2, “Discovering Enterprise Resources in Multiple vSphere Environments,”](#) on page 92.
- 10 Discover the VM images on the vCenter Server and populate the Orchestration Console Explorer tree.
- 10a From the main menu, select *Provision > Select VM Hosts and Repositories* to display the Discover VM Hosts and Repositories dialog box.
- 10b In the Discover VM Hosts and Repositories dialog box, select the `vsphere` job, then click *OK*.
-
- TIP:** Ensure that this job completes before proceeding to [Step 10c](#): Repositories where VMs might reside must be discovered prior to any attempt to discover VM images residing there.
-
- 10c From the main menu, select *Provision > Discover VM Images...* to open the Discover VM Images dialog box.
- 10d In the Source Repositories table of the Discover VM Images dialog box, 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.

Policy Configuration Summary for the vSphere Provisioning Adapter

The following table provides detailed information about other 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.

Policy Name	Description	Additional Details
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 (<code>resource.vm.pool</code>).
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 83 .
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 vsphere.policy. Make sure 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 args in the schedule that is created by default installation of the Cloud Manager Orchestration Server.	For more information, see “Configuring the vSphere Update Client” on page 88 .
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 <i>automatic</i> (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 .

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.

Use the following steps 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.

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. For more information, see [“Enabling VNC Access to vSphere 5 VM Guest Consoles”](#) in the [NetIQ Cloud Manager Component Reference](#).

Use the following steps to set up VNC session connectivity for the VM managed by the vsphere provisioning adapter job.

NOTE: 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.

- 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 box.
- 3 In the Add Credential dialog box, 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 5900-5964. If you want to provide remote capabilities to more than 65 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 `vsphere.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 `vsphere.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 box.
 - 6c** In the *Source Policies* list of the Policy Selection dialog box, 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 box.
 - 7c** In the *Source Policies* list of the Policy Selection dialog box, 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. For more information, see [“Enabling VNC Access to vSphere 5 VM Guest Consoles”](#) in the *NetIQ Cloud Manager Component 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, make sure you configure the vSphere Update Client beforehand. For more information, see [“Configuring the vSphere Update Client”](#) on page 88.

- 11 Right-click the VM object and select *Launch Remote Desktop* to open the login dialog box for the VNC session.
- 12 In the login dialog box, 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
<code>resource.vnc.ip</code>	The IP address of the VM host where the VM is running
<code>resource.vnc.port</code>	The port currently assigned to the VM. The value is -1 if VNC is disabled for the VM.
<code>resource.vnc.credential</code>	The credential containing the VNC password. This is the name of the credential itself, not the username or the password contained in the credential.
<code>resource.vnc.mode</code>	Determines how VNC port assignments are handled. This value must be <code>automatic</code> , <code>manual</code> , or <code>off</code> . <ul style="list-style-type: none">♦ If <code>mode = automatic</code>: the Orchestration Server attempts to select the next available VNC port.♦ If <code>mode = manual</code>: The port value specified in the VM's <code>resource.vnc.port</code> fact is used.♦ If <code>mode = off</code>: The VNC console is disabled.
<code>resource.remotedesktop</code>	Controls enabling or disabling the Launch Remote Desktop action in the Orchestration Console.

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 *NetIQ Cloud Manager Component Reference*.

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 the read-only cluster-related facts for these discovered clusters in the following Orchestration documentation references:

- ♦ “Orchestration Server Facts in a VM Host Residing in a Cluster” in the *NetIQ Cloud Manager Component Reference*
- ♦ “Orchestration Server Facts in the VM Host Cluster Object” in the *NetIQ Cloud Manager Component Reference*
- ♦ “Orchestration Server Facts in VMs Hosted in Clusters” in the *NetIQ Cloud Manager Component Reference*

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

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

Configuring the vSphere Update Client

To configure the vSphere Update Client:

- 1 Create a proxy user:
 - 1a In the Orchestration Console, click *Actions > Create User* to open the Create a New User dialog box.
 - 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 username must contain the word “proxy,” for example, *my_proxy*, or *proxy1*.
- 2 Modify the vSphereUpdate.policy (or modify the jobargs in the 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.
- 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.

The vSphereUpdate Monitor Job

The vSphereUpdate monitor job is located in the “all” jobs group. It is associated with both the vsphere policy (for VCenter configuration information) and the vSphereUpdate policy. The vSphereUpdate policy specifies the following cluster-related facts. You can modify these facts to accommodate your environment.

Table 11-3 Cluster-Related Facts in the vSphereUpdate Policy

Fact Name	Type	Description
jobargs.zos.proxy.user	String	<p>An administrative user used by the Orchestration Console to log in to the Orchestration Server in order to perform update operations there.</p> <p>You must create an administrative user for this purpose, if you have not already done so.</p> <p>The name of this user must contain the word “proxy,” for example, <i>my_proxy</i>, or <i>proxy1</i>. When you change the value of this fact, you must restart the Orchestration Server.</p> <p>For information about configuring the vSphere Update Client, see “Configuring the vSphere Update Client” on page 88.</p>
jobargs.zos.proxy.passwd_validity	Integer	<p>The amount of time (measured in seconds) that the <code>zos.proxy.user</code> password is valid.</p> <p>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.</p> <p>When the password expires, the Orchestration Console is automatically restarted with a new password the next time that the monitor job runs.</p>
jobargs.debug	Boolean	<p>Specifies whether you want extra verbose debug logging sent to a job log.</p> <p>NOTE: The client logs its output to the <code>log.txt</code> and <code>err.txt</code> files located in <code><agent_install_dir>/node.default/.vSphereUpdate/<hostname>/<vcenterId></code>.</p>
jobargs.verbose	Boolean	<p>Specifies whether you want verbose logging sent to a job log.</p> <p>This fact is implicitly set when <code>jobargs.debug</code> is set.</p>

Fact Name	Type	Description
jobargs.mode	String	<p>The value for this fact can be optionally set to "clear." This resets the passwd_validity and forces a restart on the next invocation where the mode is not set.</p> <p>The value can also be set to "stop" to stop all running update clients.</p>

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.

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 box.
- 4 In the Source Policies list, select *vsphere_assignPool*, click *Add* to move it to the Associated Policies list, then click *OK*.

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, follow these steps:

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

- 3b In the dialog box, 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 box.

- 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.

11.1.2 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.

- ♦ [“Creating Accounts for Each vCenter Environment” on page 92](#)
- ♦ [“Configuring the vsphere.vcenters Fact to Include All Accounts Representing a vCenter Server” on page 93](#)
- ♦ [“Optionally Specifying an Authentication Certificate for Each vCenter Server” on page 94](#)
- ♦ [“Running Discovery” on page 95](#)

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 box.
 - 3a Fill in the fields of the dialog box:
 - ♦ **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.

When 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 93](#).

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 92](#). 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 93](#)
- ♦ [“Creating a Policy to Apply to Each Orchestration Resource in the Respective vSphere Environments” on page 93](#)

When you have used one of these methods, continue with [“Optionally Specifying an Authentication Certificate for Each vCenter Server” on page 94](#).

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 92](#)) 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 box, 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 NetIQ Cloud Manager Component Reference](#).

If you want to connect multiple vCenter Servers, make sure 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 92.

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 \(http://pubs.vmware.com/vsphere-50/topic/com.vmware.wssdk.dsg.doc_50/sdk_sg_server_certificate_Appendix.6.4.html#991190\)](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 Make sure 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 43](#).
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 box.
 - 4b Fill in the fields of the dialog box:
 - ♦ **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 *NetIQ Cloud Manager Component Reference*.

4c Click *Add*.

5 Repeat [Step 4](#) for all of the vCenter Servers you want to connect to.

When you have completed the authentication configuration, continue with “[Running Discovery](#)” on [page 95](#).

Running Discovery

When the Orchestration Server is properly configured, you can use the following steps to discover the VM images on each 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 box.

2 In the Discover VM Hosts and Repositories dialog box, 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.

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

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 box, 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](#).

11.2 Configuring the Citrix XenServer Provisioning Adapter

If you manage a Citrix XenServer environment, you can use the NetIQ Cloud Manager XenServer provisioning adapter job to help you manage that environment. The `xenserv` provisioning adapter job is automatically deployed when you start the Orchestration Server.

For more information about the `xenserv` provisioning adapter policy, see “[The Citrix XenServer Provisioning Adapter](#)” in the *NetIQ Cloud Manager Component Reference*.

This section includes the following information:

- ♦ [Section 11.2.1, “Deploying the Citrix XenServer Provisioning Adapter,” on page 96](#)
- ♦ [Section 11.2.2, “Configuring the Citrix XenServer Updater,” on page 98](#)
- ♦ [Section 11.2.3, “Configuring Orchestrator for Personalization with XenServer,” on page 98](#)
- ♦ [Section 11.2.4, “Using Xen VNC Proxy to Establish a Remote Desktop Connection to XenServer VMs,” on page 98](#)

11.2.1 Deploying the Citrix XenServer Provisioning Adapter

The Citrix XenServer Provisioning Adapter uses the XenAPI management API to connect to and manage XenServer hosts. Unlike previous versions of the XenServer Provisioning Adapter, this adapter requires no additional software be installed on the XenServer host.

To configure the provisioning adapter, perform the following steps in the Cloud Manager Orchestration Console:

- 1 Create a XenServer credential for each user ID/password combination used for XenServer pool master hosts:
 - 1a Navigate to the Grid object for the Orchestrator server in the explorer panel
 - 1b Select the *Authentication* tab in the management panel
 - 1c Under “Stored Credentials”, click *Add Credential*
 - 1d Fill in the following fields:
 - Name:** The name to refer to this credential as
 - User:** Enter `root`
 - Secret:** Enter the root account’s password for the XenServer pool master
 - Type:** Select `xenserv` from the dropdown list
- 2 Create an account for each XenServer pool master host in the `xenserv` job:
 - 2a Expand the *Jobs->provisionAdapters* branch of the Grid in the explorer panel
 - 2b Select the *Job Configuration* tab in the management panel
 - 2c Under “Accounts”, click *Add*
 - 2d Fill in the following fields and save the job:
 - Account Name:** The name to use for this account
 - VM Host IP/DNS Address:** The IP or DNS address of the XenServer Pool master server.
 - Credential Name:** The name of the credential entered in [Step 1d](#)
 - Use SSL:** Check this option to encrypt the API calls over the network

NOTE: Using SSL is generally not necessary, and adds significant overhead to the traffic and processing. Only enable this option if encryption is necessary.


- 3 Associate the `xenservClient` policy with the system that will execute this job:
 - 3a Expand the *Resources* branch of the Grid in the explorer panel and select the system that will execute the job
 - 3b Select the *Policies* tab and click the *Choose* button
 - 3c Select the `xenservClient` object in the *Source Policies* list and click the *Add* button
 - 3d Click the *OK* button
- 4 Discover the XenServer hosts and repositories by clicking the *Provision* menu and selecting *Discover VM Hosts and Repositories...*
- 5 When the `Provision(xenserv)` job has completed, discover the VM images in the discovered hosts by clicking the *Provision* menu and selecting *Discover VM Images...*

NOTE: In order for VMs to be properly discovered, they must have the XenServer tools installed and the OS information populated prior to discovery.

If a single provisioning adapter is deployed within a single directly reachable network, nothing special needs to be done in order for the adapter to work. This applies both with multiple pools on the same network as well as a pool on one side of a firewall, as shown in [Figure 11-1](#).


This type of configuration is typically used when multiple networks are separated by a firewall that limits communication between the networks.

Figure 11-1 Single provisioning adapter per network



If, as shown in [Figure 11-2](#), multiple provisioning adapters are used within a single network, it is also necessary to define restrictions using the `xenservClientHostRestriction-Template` policy. These restrictions are used to limit which XenServer pool masters each provisioning adapter connects to.

Figure 11-2 Multiple provisioning adapters per network



To define these restrictions, perform the following steps:

- 1 Navigate to the `xenservClientHostRestrictions-Template` policy in the *Policies* tree.
- 2 Create a copy of the policy for the first instance of the provisioning adapter.
- 3 Modify the copy to list the account name or names that the client has access to, following the instructions in the comments in the policy.

- 4 Associate the policy with the host the provisioning adapter job is associated with.
- 5 Repeat [Steps 2-4](#) for each provisioning adapter instance.

11.2.2 Configuring the Citrix XenServer Updater

To keep the discovered facts about VMs hosted on Citrix XenServer updated, you can enable the schedule for the XenServer Updater job. To do this:

- 1 Open the Scheduler view
- 2 Select the XenServer Updater schedule
- 3 Click the *Enable* button to enable the daemon
- 4 (Optional) To start the job immediately, click the *Run Now* button

Once the daemon is started, changes to the facts tracked by Orchestrator will be reflected as they are made to the VM.

11.2.3 Configuring Orchestrator for Personalization with XenServer

When using personalization with VMs hosted on Citrix XenServer 6.0 and later, it is necessary to have a shared ISO library configured. Orchestrator will upload a customizer LiveCD image to that library that will handle all the personalization (server name, DHCP/network configuration, autoprep/sysprep and other configuration) using that ISO.

For XenServer 5.6, it is necessary to manually upload the ISO to the ISO library:

- 1 Copy the CMOS_Customizer_LiveCD.i686-x.y.z.iso file from `/opt/novell/zenworks/zos/server/doc/install` directory into the shared ISO library
- 2 Rescan the ISO library in XenCenter to verify the upload has completed.
- 3 In the Orchestration Console, browse to the ISO library, right click, and select *Discover Disks*

Once these steps are completed, personalization of VMs on the XenServer host will run as expected.

11.2.4 Using Xen VNC Proxy to Establish a Remote Desktop Connection to XenServer VMs

NetIQ Cloud Manager uses the Xen VNC proxy (xvp) server to provide a password-based connection to all of the guest VM consoles that are hosted on a single Citrix XenServer that is connected to a single Cloud Manager Orchestration Server.

This section includes information about how to install and configure xvp for use with Cloud Manager.

- ♦ [“Installing the Xen VNC Proxy Packages” on page 99](#)
- ♦ [“Configuring Xvp Credentials in the Orchestration Server for the Citrix XenServer Environment” on page 99](#)
- ♦ [“Understanding How Xen VNC Proxy Works in the Orchestration Environment” on page 100](#)
- ♦ [“Cloud Manager Console Actions on XenServer VMs Configured To Use Xen VNC Proxy” on page 101](#)
- ♦ [“Known Issues with Xen VNC Proxy Remote Console Usage” on page 101](#)

Installing the Xen VNC Proxy Packages

If you want a supported method of launching a remote console of a VM managed by Citrix XenServer, you need to install the xvp package provided by NetIQ on the Cloud Manager installation ISO.

To install xvp packages:

- 1 Mount the Cloud Manager installation ISO on a network computer running a supported version of SUSE Linux Enterprise Server (SLES). This computer should not be part of the existing Citrix Xen environment. It must also have the NetIQ Cloud Manager Orchestration Agent installed on it.
- 2 On the computer where you are installing xvp, start YaST and select *Software Management*.
- 3 In the YaST *Software Management* view, select the *Xen VNC Proxy* install pattern, then click *Accept* to install the packages. The pattern includes two xvp packages:
 - ♦ libxenserver
 - ♦ xvp
- 4 Start the Orchestration Agent on the SLES computer where you installed the xvp proxy.
- 5 In the Orchestration Console toolbar, select *Resources*, select the registration icon to open the Resource Registration Monitor dialog box, then click *Accept > OK* to register the new resource.

When the resource is registered, it is automatically discovered as an xvp host.

Configuring Xvp Credentials in the Orchestration Server for the Citrix XenServer Environment

After the host discovery, you need to set up the credentials that allow xvp to open ports for the VNC sessions to the Citrix XenServer VMs.

- ♦ [“Creating Credentials for Individual Citrix XenServer VMs” on page 99](#)

Creating Credentials for Individual Citrix XenServer VMs

VMs are most commonly provisioned and given VNC credentials by Cloud Manager business owners, who own and control those VMs as managed workloads. These are saved in the Orchestration Server credential store.

If you want to create VNC credentials manually for Citrix XenServer VMs you manage with Cloud Manager Orchestration Console, you can use the following steps:

- 1 In the Explorer tree of the Orchestration Console, select the Grid object for the Orchestration Server that communicates with the Citrix Xen environment, then in the Admin view, select *Authentication* to open the Authentication page.
- 2 In the Stored Credentials subpanel, select *Add Credential* to open the Add Credential dialog box.
- 3 In the Add Credential dialog box, fill in the fields to create a new VNC credential set for a VM. computer.
 - ♦ **Name:** This is a required field. Provide a name that you want to use to identify this credential set.
 - ♦ **User:** This is a required field. Enter a username you want use in the VNC session for this VM.
 - ♦ **Secret:** This is a required field. Enter a password you want use in the VNC session for this VM.

IMPORTANT: This password must be no more than 10 characters. Passwords with more than 10 characters do not store properly.

- ♦ **Type:** Select *VNC* as the credential type.
- 4 Click *Add* to save the credential information.
- 5 In the Explorer tree, select a VM that is hosted by the Citrix XenServer computer “host”.
- 6 Apply the newly created VNC credential for this VM.
 - 6a Select the Info/Facts tab to open the Info/Groups page for this VM.
 - 6b On the Info/Groups page, scroll to the *VNC Credential* field in the *Resource Information* subpanel.
 - 6c In the VNC Credential field, open the drop-down menu to list the configured credentials, then select the name of the credential that you created in [Step 3](#).
 - 6d Click *Save* to commit the change.
 - 6e In the Explorer tree, right-click the VM to which you just added a credential, then click *Apply Config* or *Save config* to enable the credential.

This action populates the following facts on the VM:

- ♦ resource.vnc.port
- ♦ resource.vnc.ip

Understanding How Xen VNC Proxy Works in the Orchestration Environment

When the host discovery runs on the SLES resource where the Orchestrate Agent is running, it checks for the xvp service at `/etc/init.d/xvp`. If the service is present, four facts are created. The following table lists these facts, their values and purpose.

Table 11-4 Facts for the Xvp Service Listed in the SLES Resource

XVP Fact Name	Default Port Value	Purpose
resource.xvp.beginport	6901	<ul style="list-style-type: none">♦ User configurable.♦ Instructs the xvp computer which port to start using.
resource.xvp.freeport	6901	<ul style="list-style-type: none">♦ User configurable.♦ Informs the xvp machine which port to use for the next provisioned VM♦ Value increments automatically.♦ Ports assigned to destroyed VMs are stored in <code>var/xenservXVP_freeports.txt</code>, to be reused later.
resource.xvp.vncportrange	100	<ul style="list-style-type: none">♦ User configurable.♦ Provides information for the xvp computer regarding how many ports it should use for proxy connections.
resource.xvpHost	true	<ul style="list-style-type: none">♦ Identifies the system as an xvp proxy server.

Cloud Manager Console Actions on XenServer VMs Configured To Use Xen VNC Proxy

The following table lists some of the actions you can perform on a Citrix XenServer VM that you have configured with xvp credentials for using a remote console session.

VM Action	Result
Apply Config or Save Config	The Orchestration Server makes an entry in the xvp configuration file for the selected VM.
Migrate or Move	The Orchestration Server moves the “VM configuration” information in the XVP configuration file to the XVP configuration file of the destination XenServer.
Destroy	The Orchestration Server deletes the “VM configuration” information in the xvp configuration file. If there are no such entries in the file, the server deletes the entire file, along with the corresponding entry from the main configuration file (<code>/etc/xvp.conf</code>).

Known Issues with Xen VNC Proxy Remote Console Usage

There are some known issues with Cloud Manager 2.1 remote console connections to Citrix XenServer VMs via xvp:

- Only one xvp proxy server can be registered on the Orchestration Server. If you determine that network traffic becomes too much for this single proxy to efficiently handle its remote connections, you can deploy another xvp, but you also need to deploy an additional Orchestration Server to manage it.
- Occasionally, a workload (that is, a Citrix XenServer VM) provisioned from the Cloud Manager Application Server Console fails to properly configure the `resource.vnc.ip` fact and the `resource.vnc.port` fact. Use the *Apply Config* action on the VM in the Orchestration Console to correct the configuration of these facts.

11.3 Configuring the Hyper-V Provisioning Adapter

The hyperv provisioning adapter job deploys the `hyperv.policy` with the job. This policy contains the facts and constraints that the hyperv provisioning adapter job uses for checking whether the Hyper-V server host is registered to the Orchestration Server, and whether that host is up and running. By default, the optimal values are preset for the configuration of the job and joblets in the policy. We strongly recommend that you do not edit this policy.

The following additional configuration information is included in this section:

- [Section 11.3.1, “Ensuring that the Orchestration Server Discovers Hyper-V VMs,” on page 102](#)
- [Section 11.3.2, “Configuring the Provisioning Adapter to Discover iSCSI Target Repositories,” on page 102](#)
- [Section 11.3.3, “Configuring the Provisioning Adapter for Sysprep,” on page 102](#)
- [Section 11.3.4, “Enabling a Remote Console Session for a Hyper-V VM,” on page 102](#)
- [Section 11.3.5, “Configuring Hyper-V Linux VMs to Enable Visibility of Added vDisks,” on page 103](#)

For information about troubleshooting the hyperv provisioning adapter job, see [“Troubleshooting Hyper-V VM Provisioning Operations”](#) in the *NetIQ Cloud Manager Component Reference*.

11.3.1 Ensuring that the Orchestration Server Discovers Hyper-V VMs

If you create a VM in your Hyper-V environment, but the path to that VM was not configured as the default path in the Hyper-V Manager, the Orchestration Server does not discover the VM until you edit the preferred path for the discovered repository where the VM resides. You can also create a new repository in the Orchestration Console with the preferred path to the Hyper-V VM.

11.3.2 Configuring the Provisioning Adapter to Discover iSCSI Target Repositories

If you are managing Windows VMs in a Hyper-V environment (clustered or non-clustered), the hyperv provisioning adapter must be configured to discover iSCSI target repositories in that environment if the VM is in a location other than `C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks`.

To configure the provisioning adapter for this use case:

- 1 In the Explorer tree of the Orchestration Console, select the *Repositories* group to expand the list of Repository Objects, select *hyperv*, then select the storage object associated to the Hyper-V cluster to open the admin view.
- 2 In the Info/Groups page of the admin view, find the *Preferred Storage Path* field (the `repository.preferredpath` fact).
- 3 In the *Preferred Storage Path* field, change the value to the path where the VM resides. Remember that this field considers the information in the *Root Location* field (that is `repository.location`).

This is the location where the Orchestration Server searches for VM files for use in cloning and moving. Generally, it is a path like this:

```
C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks
```

- 4 Click the *Save* icon to save the new configuration.

NOTE: If your Hyper-V environment is a Cluster Storage Volumes (CSV) environment, the VMs on the CSVs are automatically discovered by the hyperv provisioning adapter as separate repositories. Executing the *Discover VM Images* action on these repositories discovers the VMs residing there.

11.3.3 Configuring the Provisioning Adapter for Sysprep

As with other VMs provisioned by the Orchestration Server, sysprep does not work on Hyper-V 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 114](#).

11.3.4 Enabling a Remote Console Session for a Hyper-V VM

If you invoke the VNC console for a Hyper-V VM (referred to as a “workload”) from the Cloud Manager Web Client, the VNC console does not launch.

Installing the Orchestration Agent on the VM and executing the *Apply Config* action lets you launch a VNC session from Cloud Manager to the Hyper-V “workload” desktop.

To install the agent to the VM:

- 1 In Explorer tree of the Orchestration Console, select the VM that you want to observe in a remote session, then right-click and select *Shutdown*.
- 2 Right-click the now idle VM, then select *Install Agent*.
- 3 Right-click the VM, then select *Start*.
- 4 When the VM appears online again in the list of resources, right-click the VM again and select *Apply Config*.

11.3.5 Configuring Hyper-V Linux VMs to Enable Visibility of Added vDisks

If you plan to add an additional vDisk to the Hyper-V Linux VM at some point, you need to further configure the VM so that the vDisk is visible. To do this, you need to install Microsoft Linux Integration Components for Linux. See the [Microsoft download site \(http://download.microsoft.com/download/4/2/7/4273D9CF-3FC3-4A91-8204-9E0D4DE2027C/Linux%20Integration%20Components%20Read%20Me.pdf\)](http://download.microsoft.com/download/4/2/7/4273D9CF-3FC3-4A91-8204-9E0D4DE2027C/Linux%20Integration%20Components%20Read%20Me.pdf) for more information.

11.4 Configuring the SUSE Xen Provisioning Adapter

The xen provisioning adapter job has prepackaged policies that are deployed with the job. These policies run when the job is deployed and manage the Xen hosts and VMs in the grid. The policy settings are applied to all the VMs in the grid.

For more information about the provisioning adapter policies, see “[Provisioning Actions Supported by the SUSE Xen Provisioning Adapter](#)” in the *NetIQ Cloud Manager Component Reference*.

For information about troubleshooting the xen provisioning adapter job, see “[Troubleshooting SUSE Xen VM Provisioning Actions](#)” in the *NetIQ Cloud Manager Component Reference*.

11.4.1 Cloud Manager Orchestration Defaults in a SUSE Xen Cluster

The following content describes certain Cloud Manager Orchestration Server installation and configuration defaults when it is installed in a SUSE Xen clustering environment.

- ♦ “[A New Orchestration Public JDL Library](#)” on page 103
- ♦ “[The xen Provisioning Adapter](#)” on page 104
- ♦ “[Changes in the xendConfig Job](#)” on page 106
- ♦ “[The Xen.CMOS OCF Script](#)” on page 106

A New Orchestration Public JDL Library

A new Public JDL Library (`linuxha.pylib`) provides the common APIs necessary for Orchestration to interact with the SLES 11 HAE clustering stack (that is, Pacemaker).

The xen Provisioning Adapter

Some minor exceptions to the behavior of SUSE Xen VMs managed by the Orchestration Server in a Pacemaker cluster include the following:

- ♦ Building VMs (that is, using the VM Builder through the Orchestration Server *Build* action) is not supported in a Xen cluster. You first need to build the VMs in a non-clustered environment and then afterwards provision them to the Xen cluster.
- ♦ The *Apply Config* action is not supported for VMs running within a Xen cluster.
- ♦ The *Pause*, *Resume*, *Suspend*, and *Launch Remote Desktop* (VNC console) actions are supported in the cluster only when the VM is configured to use the [Xen.CMOS script](#). These actions are unavailable if you use the default Xen OCF script provided by the SLES 11 HAE product.
- ♦ During the *Discover VM Hosts and Repositories* action, the xen provisioning adapter discovers any configured Xen clusters and models them as follows:
 - ♦ Creates a `VmHostCluster` object using the name of the clustered Orchestration Agent by default.
 - ♦ Discovers all nodes associated with the create `VmHostCluster` operation and models them as regular VM host Grid objects with the `vmhost.cluster` fact specified (this identifies the VM host as a member of a `VmHostCluster`).
 - ♦ Creates repositories based on file systems under the control of the clustering stack. These are resources that define the `Filesystem` OCF script in the cluster CIB. The provisioning adapter does not discover any shared storage that is not under control of the cluster stack. You must manually add this storage, if necessary.
- ♦ When VMs are discovered (that is, when the *Discover VM Hosts and Repositories* action is executed), they are configured with a set of custom facts. These facts are stored on the VM Grid object. They are listed and described in the table below.

Custom Fact Name	Description
<code>resource.vm.linuxha.cib_xml.id</code>	The unique ID of this VM as known by the cluster CIB. The fact is used as an identifier when performing actions on a VM running within a Xen Pacemaker cluster.
<code>resource.vm.linuxha.cib_xml</code>	<p>The stored CIB XML definition for this VM in the cluster. When a VM is discovered (that is, when the <i>Resync State</i>, <i>Check Status</i>, or <i>Discover VM Hosts and Repositories</i> actions are executed) this fact is overwritten with new data queried from the cluster CIB.</p> <p>For all other actions (for example, <i>Provision</i>, <i>Save Config</i>, etc.), the cluster's CIB definition for the VM is replaced with the contents of this fact.</p>

- ♦ The new `xenClusterVmDefaults` policy is associated by default to the `VMs_xen` Resource Group. The policy defines a set of default facts on a VM grid object that belongs to a Xen cluster. The facts are listed and described in the table below.

Fact Name	Description
<code>resource.vm.linuxha.cib_xml_default</code>	The default CIB primitive to be used in the case where a CIB definition for this VM does not exist (that is, the VM is not defined in the cluster CIB and the value for the <code>resource.vm.linuxha.cib_xml</code> fact is empty).

Fact Name	Description
<code>resource.vm.linuxha.clear_migrate_constraints_on_shutdown</code>	<p>Used when shutting down (or restarting) the vm. If set, any location constraints created by migrate will be cleared at shutdown time.</p> <p>If there are specific location constraints created by an administrator, they are not cleared (unless they follow the naming convention for standard migrate constraints).</p>
<code>resource.vm.linuxha.failcount.threshold</code>	<p>Defines the allowed threshold for the failcount value used during the provisioning of a VM to a cluster. If the specified failcount threshold has been reached on all nodes in the cluster, the failcount will be reset back to 0 before attempting the start of the VM. This value can be one of: '+INFINITY', or '<Integer>'. To disable this feature, specify a value of '-1'.</p>
<code>resource.vm.linuxha.migrate.timeout</code>	<p>Timeout (in seconds) to wait for a clustered VM to reach the running state following migration. A value of 0 means wait forever. If set, this value overrides <code>vm.linuxha.vm_state.wait.timeout</code> (for migrate).</p>
<code>resource.vm.linuxha.provision.timeout</code>	<p>Timeout (in seconds) to wait for a clustered VM to reach the running state. A value of 0 means wait forever. If set, this value overrides <code>vm.linuxha.vm_state.wait.timeout</code>, and is set as the start operation timeout.</p>
<code>resource.vm.linuxha.shutdown.timeout</code>	<p>Timeout (in seconds) to wait for a clustered VM to reach the running state. A value of 0 means wait forever. If set, this value overrides <code>vm.linuxha.vm_state.wait.timeout</code>, and is set as the stop operation timeout and the VM's <code>shutdown_timeout</code> in the CIB.</p> <p>To use per-VM values, set the value of this fact to the empty string (''). The default is to use the value of the <code>resource.vm.shutdown.timeout</code> fact.</p>
<code>resource.vm.linuxha.vm_state.wait.interval</code>	<p>Defines the interval (in seconds) when the desired VM state should be re-checked (sleep interval). The default value is 5 seconds.</p>
<code>resource.vm.linuxha.vm_state.wait.timeout</code>	<p>Timeout (in seconds) to wait for a desired VM state. A value of -1 means wait forever. The default value is 120 seconds (2 min).</p>

- ♦ The new `xenClusterResource` policy is associated by default to the `Clusters_xen` Resource Group. The policy defines a set of default facts on a Xen cluster grid object. The facts are listed and described in the table below.

Fact Name	Description
<code>resource.linuxha.joblets.default_slots</code>	Specifies the default number of joblets slots to create for a Xen Cluster resource.
<code>resource.linuxha.cibadmin.cache.enabled</code>	If the value for this fact is true, the cluster CIB is cached when invoking queries. If the value is false, the CIB is queried directly.
<code>resource.linuxha.cibadmin.cache.lifetime</code>	The amount of time (in seconds) for which the cached CIB should stay valid.
<code>resource.linuxha.cibadmin.cache.invalidate_on_write</code>	Specifies (true or false) whether the CIB cache should be invalidated when writing to the CIB.

Changes in the xendConfig Job

Generally, the `xendConfig` job configures each SUSE Xen VM host (that is, it modifies `/etc/xen/xend-config.sxp`). The configuration includes but is not limited to the following functions:

- ♦ VNC console sessions to a VM ('vnc-listen', 'vnc-passwd')
- ♦ VM migration between Xen VM hosts ('xend-relocation-server', 'xend-relocation-address', 'xend-relocation-port', 'xend-relocation-hosts-allow').

Currently, when the Orchestration Agent is running in a clustered environment (that is, it is active on only one host in the cluster at a time), it is not possible to use the `xendConfig` job to configure all of the VM hosts contained within the Xen cluster using the `xendConfig` job because the agent does not have access to all hosts.

If the `xendConfig` job detects that it is running on a clustered VM host, the following message (or similar) is displayed:

```
WARNING: cluster-xen2_xen is a member of a XEN cluster. Xend should be manually
configured on this host. Skipping..
```

To work around this issue, manually configure the `/etc/xen/xend-config.sxp` file on every cluster node as appropriate, then restart the xen service (`rcxend restart`).

The Xen.CMOS OCF Script

The Cloud Manager Orchestration Server extends the SUSE Xen OCF script provided with the SLES 11 HAE product. This enhanced Xen OCF script, called `Xen.CMOS`, resides in `/usr/lib/ocf/resource.d/heartbeat/`.

By modifying the `resource.vm.linuxha.cib_xml` fact or the `resource.vm.linuxha.cib_xml_defaults` fact, you can define any or all Xen VMs to use this enhanced script (the default configuration, which provides VNC session functionality and the ability to execute the *Pause*, *Resume*, and *Suspend* actions on those VMs) or you can configure them to run without this functionality by using the original Xen OCF script.

For example, in order to revert to unenhanced Xen functionality, you would change the CIB XML text from this:

```
<primitive class="ocf" id="sles10-pv" provider="heartbeat" type="Xen.CMOS">
```

to this:

```
<primitive class="ocf" id="sles10-pv" provider="heartbeat" type="Xen">
```

11.5 Configuring the KVM Provisioning Adapter

If you manage a KVM (kernel based virtual machine) hypervisor environment where the hypervisor runs on a SUSE Linux Enterprise Server (SLES) 11 SP1 host machine, you can use the Orchestration KVM provisioning adapter to help manage its VMs and to expose guest OS machines to Cloud Manager.

The KVM provisioning adapter uses job configuration settings specified on the kvm job, which is automatically deployed when you start the Orchestration Server. You can select these settings when you select the kvm job and open the Job Configuration tab in the job admin view of the Orchestration Console.

Figure 11-3 The Job Configuration Page of the KVM Provisioning Adapter Job

The page includes two subpanels:

- ♦ [Authentication Settings](#)
- ♦ [Debug Settings](#)

There are additional `kvm.policy` file that are used by the provisioning adapter. You should retain the default values for these settings.

For more information about the how the KVM provisioning adapter works and for a list of the facts in the kvm provisioning adapter policy, see [“The KVM Provisioning Adapter”](#) in the *NetIQ Cloud Manager Component Reference*.

11.5.1 Authentication Settings

These settings relate to the type of authentication used between KVM hosts in your hypervisor’s environment. You need to set up some form of authentication between KVM hosts so that migrate operations function correctly. The Orchestration Server also uses the authenticated channel to

optimize how VM registrations are moved between hosts, though it falls back to a less-efficient method of running joblets on both resources, if necessary. The Authentication Settings subpanel includes the following fact settings:

- ♦ **Authentication Type:** The default type of authentication used by the libvirt libraries. This setting can be overridden on a per-host basis by setting the `vmhost.libvirt.authentication.type` fact. Valid values are `ssh` or `tls`.

The setting is listed in the Fact Editor as a String type, named `job.libvirt.authentication.type.default`.

- ♦ **Default User:** The default user to use to connect over a secure channel to the libvirt user on a given host. This setting can be overridden on a per-host basis by setting the `vmhost.libvirt.authentication.user` fact on any given host.

The setting is listed in the Fact Editor as a String type, named `job.libvirt.authentication.user.default`.

You are responsible to set up the authentication channel and to advertise it to the Cloud Manager Orchestration Server (Cloud Manager does not attempt to automatically create these connections if they do not exist). For setup instructions for libvirt authentication on SUSE-based hosts, see [Chapter 7, “Connecting and Authorizing”](http://www.novell.com/documentation/sles11/book_sles_kvm/data/cha_libvirt_connect.html) (http://www.novell.com/documentation/sles11/book_sles_kvm/data/cha_libvirt_connect.html) in the *SLES 11 Virtualization with KVM Administration Guide*.

The SSHConfigure job can help you set up SSH authentication, but you must run it manually. The job:

- ♦ Sets up ssh keys for the root user on each kvm host.
- ♦ Imports public keys into the Credential Manager from each host.
- ♦ Propagates hosts’ public keys to other hosts and makes them available for use.

Use the following steps to run the SSHConfigure job manually:

- 1 Create a Resource Group that includes each kvm host you want to configure for SSH.
- 2 In the Job Configuration page of the Orchestration Console admin view for the SSHConfigure job, edit the settings.
 - 2a Enter the desired group name for SSH keys. This is the group name for keys in the Credential Manager.
 - 2b Define the desired SSH user name.
 - 2c Enter the name of the group you created in [Step 2a](#) above into the *Desired Host Group* field.
- 3 Run the SSHConfigure job by using the SSHConfigure schedule.

11.5.2 Debug Settings

These settings control the amount of information that is reported to a job log. The Debug Settings subpanel includes the following fact settings:

- ♦ **Log level:** Specifies the level of verbosity for the information recorded in the job log. Values of “OFF”, “CRITICAL”, “ERROR”, “WARN”, “INFO”, “DEBUG”, “FINE”, “FINER”, and “FINEST” produce increasing detail.

The setting is listed in the Fact Editor as a String type, named `job.debugLevel`.

- ♦ **Enable tracing:** Select this check box to log a TRACE statement to the job log each time an annotated job method is entered or exited. The only reason to enable this is to provide additional detail when reporting a bug.

The setting is listed in the Fact Editor as a Boolean type, named `job.traceEnabled`.

- ♦ **Trace width:** Trace statements are formatted so that they do not exceed this line length. The default value is 120.

The setting is listed in the Fact Editor as an Integer type, named `job.traceWidth`.

- ♦ **Trace depth:** Specifies how much argument and return value data is displayed on trace statements. A value of 0 (zero) suppresses method argument and return values from the job log. Larger values indicate how many levels of data should be logged if arguments are nested data structures.

The setting is listed in the Fact Editor as an Integer type, named `job.traceDepth`.

12 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 12.1, “Understanding and Configuring Sysprep,” on page 111](#)
- ♦ [Section 12.2, “Understanding and Configuring Autoprep,” on page 125](#)

12.1 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, make sure that you have performed the following prerequisite tasks:

- ♦ Create a Windows VM by using the hypervisor.
- ♦ Make sure the VM is discovered by the Orchestration Server.
- ♦ Configure the VM (or a template of the VM) by using the information in [Section 12.1.2, “Setting Sysprep Facts in the Orchestration Console,” on page 112](#) and in [Section 12.1.3, “Using the Sysprep deploy.cab Files,” on page 120](#).

When the prerequisites are met, you can proceed to sysprep the VM by using the information in [Section 12.1.4, “Applying Sysprep Facts,” on page 122](#).

This section includes the following information:

- ♦ [Section 12.1.1, “How Sysprep Works,” on page 112](#)
- ♦ [Section 12.1.2, “Setting Sysprep Facts in the Orchestration Console,” on page 112](#)
- ♦ [Section 12.1.3, “Using the Sysprep deploy.cab Files,” on page 120](#)
- ♦ [Section 12.1.4, “Applying Sysprep Facts,” on page 122](#)
- ♦ [Section 12.1.5, “Example Sysprep Scenarios,” on page 123](#)
- ♦ [Section 12.1.6, “Known Sysprep Limitations,” on page 123](#)

12.1.1 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.

12.1.2 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 12-1 The Sysprep Sections of the Info/Groups Page of a VM Template Object

The screenshot displays two subpanels for configuring Windows VM sysprep facts. The top subpanel, titled "Windows Sysprep Config", contains a list of 20 configuration items, each with a text input field (currently showing "< undefined >") and a "Define" button. The items are: Change SID, Delete Accounts, Admin Password, Admin Password Plaintext, Timezone, Autologon, Autologon Count, Fullname, Org Name, Computer Name, Product ID, Run Once Command, Workgroup, Domain, Domain Admin, Domain Admin Password, Domain Admin Password Plaintext, License File Automode, and License File Autousers. The bottom subpanel, titled "Network Sysprep Config", contains three items: DNS Server IP Addresses, DNS Suffixes, and Gateway IP Addresses, each with a text input field and a "Define" button.

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 113](#)
- ♦ [“Configuring vNIC Sysprep Facts” on page 120](#)

Sysprep Facts

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

Table 12-1 Required or Optional Facts for Sysprep Configuration

String in Orchestration Console UI	Fact Name	Required/Optional	Description and Information
<i>Change SID</i>	<code><fact name="resource.provisioner.autoprep.options.changeSID" value="false" type="Boolean" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>The Windows Security ID. If this is marked as true, sysprep generates a new Security ID.</p> <p>If this is not selected, the Orchestration Server defaults the value to true, meaning a new SID is to be generated during sysprep.</p> <p>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.</p>

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
<i>Delete Accounts</i>	<code><fact name="resource.provisioner .autoprep.options.deleteAc counts" value="false" type="Boolean" /></code>	Optional ²	<p>(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.</p> <p>(Xen) This field is deprecated for Xen sysprep. Instead, ensure that the VM image has the required set of accounts from the beginning.</p> <p>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.</p>
<i>Admin Password</i>	<code><fact name="resource.provisioner .autoprep.sysprep.GuiUnatt ended.AdminPassword.value" value="" type="String" /></code>	Required ³	<p>The Admin password for this VM.</p> <p>NOTE: Only a plaintext admin password is currently supported. You should leave this field blank if <i>Admin Password Plaintext</i> is not selected.</p> <p>This fact must be specified or the personalization fails. Windows sysprep requires that this be specified.</p>
<i>Admin Password Plaintext</i>	<code><fact name="resource.provisioner .autoprep.sysprep.GuiUnatt ended.AdminPassword.plainT ext" value="false" type="Boolean" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>When this check box is selected (it has a value of true) the <i>Admin Password</i> value is entered in plain text.</p> <p>If not set, this fact defaults to true, indicating that the AdminPassword fact is a plain text password.</p> <p>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 <code>sysprep.inf</code>, then enter it as <code>AdminPassword.value</code> with this flag set to false.</p>

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
<i>Timezone</i>	<code><fact name="resource.provisioner .autoprep.sysprep.GuiUnatt ended.TimeZone" value="" type="String" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>The time zone of the new VM. For sysprep on Windows versions prior to Vista, (that is, versions using <code>sysprep.inf</code>), 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)).</p> <p>NOTE: Make sure that you use the exact text string listed under the Time Zone column heading in the table included in this Microsoft article.</p> <p>For sysprep on Windows Vista and later (that is, versions using <code>unattend.xml</code>), full string time zone names are used. Refer to the Microsoft sysprep documentation for the relevant Windows version.</p> <p>If you do not set this fact, the default time zone for <code>sysprep.inf</code> (UTC or 85) is used.</p>
<i>Autologon</i>	<code><fact name="resource.provisioner .autoprep.sysprep.GuiUnatt ended.AutoLogon" value="false" type="Boolean" /></code>	Optional ²	<p>When this check box is selected (it has a value of true) the VM automatically logs on to the Administrator account by using <i>AdminPassword</i>.</p> <p>If the value is false, logon is prompted.</p> <p>If no value is provided, the fact is set to false.</p>
<i>Autologon Count</i>	<code><fact name="resource.provisioner .autoprep.sysprep.GuiUnatt ended.AutoLogon" value="" type="Boolean" /></code>	Optional ²	<p>The limit count for the VM to automatically log on with the Administrator account. <i>AutoLogon</i> must be true for this value to be accepted.</p> <p>If a value is not specified for this fact, but <i>Autologon</i> is set to true, the value defaults to 1.</p>

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
<i>Fullname</i>	<code><fact name="resource.provisioner .autoprep.sysprep.UserData .FullName" value="" type="String" /></code>	Automatic default value provided by the Orchestration Server ¹	The user's full name required by the Windows OS installer during installation.
<i>Org Name</i>	<code><fact name="resource.provisioner .autoprep.sysprep.UserData .OrgName" value="" type="String" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>The organization name required by the Windows OS installer during installation.</p> <p>This fact is required by sysprep. If the value is not specified, the provided default is Organization Name.</p> <p>This fact is nonessential for Windows functionality.</p>
<i>Computer Name</i>	<code><fact name="resource.provisioner .autoprep.sysprep.UserData .ComputerName" value="" type="String" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>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.</p> <p>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.</p> <p>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.</p> <p>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:</p> <pre>whitespace ` ! @ # \$ ^ & * () + = [] { } \ ; : ' " , < > / ?</pre> <p>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.</p>

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
<i>Product ID</i>	<pre><fact name="resource.provisioner .autoprep.sysprep.UserData .ProductID" value="" type="String" /></pre>	Effectively required ⁴	<p>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.</p> <p>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.</p> <p>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).</p>
<i>Run Once Command</i>	<pre><fact name="resource.provisioner .autoprep.sysprep.GuiRunOn ce.Command" value="" type="String" /></pre>	Optional ²	<p>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.</p> <p>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.</p>
<i>Workgroup</i>	<pre><fact name="resource.provisioner .autoprep.sysprep.Identifi cation.JoinWorkgroup" value="" type="String" /></pre>	Automatic default value provided by the Orchestration Server ¹	<p>The Windows workgroup name. If the VM is joining a domain, use JoinDomain.</p> <p>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.</p> <p>If no domain is being joined, and this fact is not specified, the default value becomes WORKGROUP (default with Windows).</p>

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
<i>Domain</i>	<code><fact name="resource.provisioner .autoprep.sysprep.Identifi cation.JoinDomain" value="" type="String" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>The Windows domain name. If the VM is joining a workgroup, use <code>JoinWorkgroup</code>. For joining a domain, <code>DomainAdmin</code> and <code>DomainAdminPassword</code> must be defined.</p> <p>No default value is provided if this value is not set. Instead, a workgroup is joined with the default name <code>WORKGROUP</code> if no <code>Workgroup</code> fact was set. See also: Domain Admin Password (required if a value is set for this fact).</p>
<i>Domain Admin</i>	<code><fact name="resource.provisioner .autoprep.sysprep.Identifi cation.DomainAdmin" value="" type="String" /></code>	Required ³	<p>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.</p> <p>If the Domain fact does not have a value, this fact is ignored.</p>
<i>Domain Admin Password</i>	<code><fact name="resource.provisioner .autoprep.sysprep.Identifi cation.DomainAdminPassword .value" value="" type="String" /></code>	Required ³	<p>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.</p> <p>If the Domain fact does not have a value, this fact is ignored.</p>
<i>Domain Admin Password Plaintext</i>	<code><fact name="resource.provisioner .autoprep.sysprep.Identifi cation.DomainAdminPassword .plainText" value="false" type="Boolean" /></code>	Automatic default value provided by the Orchestration Server ¹	<p>When this check box is selected (it has a value of true), <code>DomainAdminPassword</code> is in plaintext.</p> <p>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 <code>sysprep.inf</code> or <code>unattend.xml</code>.</p>

String in Orchestration Console UI	Fact Name	Required/ Optional	Description and Information
<i>License File Automode</i>	<pre><fact name="resource.provisioner .autoprep.sysprep.LicenseF ilePrintData.AutoMode" value="" type="String" /></pre>	Automatic default value provided by the Orchestration Server ¹	<p>The value in this field is either PerServer or PerSeat. If it is set to PerServer, the AutoUsers fact must be set.</p> <p>A value for this fact is required for sysprep on Windows Server products. The provided default is PerServer.</p>
<i>License File Autousers</i>	<pre><fact name="resource.provisioner .autoprep.sysprep.LicenseF ilePrintData.AutoUsers" value="200" type="Integer" /></pre>	Automatic default value provided by the Orchestration Server ¹	<p>Specify value between 1 and 9999, representing the number of client licenses per seat.</p> <p>A value for this fact is required for sysprep on Windows Server products. The provided default is 5.</p>
Available in Facts tab only (not a string)	<pre><fact name="resource.provisioner .autoprep.autoprep_comple te" value="false" type="Boolean" /></pre>	Automatic default value provided by the Orchestration Server ¹	<p>Under normal circumstances, you should not need to change the value of this fact.</p> <p>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.</p> <p>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 Server 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).</p> <p>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.</p>

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 127](#).

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

12.1.3 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, Windows XP, or any previous version. 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 2003 VMs, Windows XP VMs, or earlier versions of Windows VMs. The instructions include the following:

- ♦ [“.cab File Installation Locations” on page 120](#)
- ♦ [“Detailed Instructions for Downloading .cab Files From Microsoft” on page 121](#)
- ♦ [“Detailed instructions for Copying `deploy.cab` from the Windows Installation CD or DVD” on page 122](#)

.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/windowsxp/x86/deploy.cab
dataGrid/files/sysprep/windowsxpx64/x86_64/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 for Xen and Hyper-V VMs.

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

- 1 (Conditional) Download the Windows 2003 `.exe` file containing `deploy.cab`.
 - 1a 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 2003 sysprep tool, `WindowsServer2003-KB926028-v2-x86-ENU.exe`.
 - 1b 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.
 - 1c Navigate to the extracted directory where `deploy.cab` is located.
 - 1d Copy `deploy.cab` to the [appropriate locations](#) on the Orchestration Server:
- 2 (Conditional) Download the Windows XP `.cab` file:
 - 2a From a Web browser, navigate to the Microsoft Download Center page entitled *Windows XP Service Pack 2 Deployment Tools* (<http://www.microsoft.com/downloads/en/details.aspx?FamilyId=3E90DC91-AC56-4665-949B-BEDA3080E0F6&displaylang=en>), then download the Windows XP sysprep tool, `WindowsXP-KB838080-SP2-DeployTools-ENU.cab`.
 - 2b Run the `.cab` file on a Windows physical or virtual machine to extract `deploy.cab`.

2c Navigate to the extracted directory where `deploy.cab` is located.

2d 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. Make sure 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 120](#)). For example, if your CD is the x86_64 version of Windows 2003 Server SP2, 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 121](#)) causes problems with sysprep on your VM images, using the method of copying `deploy.cab` described in this section might correct compatibility problems.

12.1.4 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 *NetIQ 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 make sure 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.

12.1.5 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 [NetIQ Cloud Manager Component 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 box, 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 [NetIQ Cloud Manager Procedures Guide](#).

12.1.6 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, make sure 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.

- If a Hyper-V VM is running during the discover process, it fails to discover the `resource.os.family` fact. This prevents the Orchestration Console from displaying the *Sysprep* and *Autoprep* options section on the *Info/Groups* tab for the VM.

NOTE: If the VM not running when the discovery occurs, the hyperv provisioning adapter discovers `resource.os.family` itself.

If you create a template from this VM, the `resource.os.family` fact is discovered and populated on the VM template admin view.

To display the sysprep/autoprep settings on a Hyper-V VM use the following steps:

1. Shut down the Hyper-V VM that has the problem.
 2. In the Explorer Tree, right-click the VM you have shut down, then select *Resync State*.
 3. In the Orchestration Console, Shift+click the Refresh  icon or restart the Orchestration Console to refresh all of the objects and their facts in the Resource admin view.
- ♦ 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 114](#).

- ♦ 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.

12.2 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 12.2.1, “How Autoprep Works,” on page 125](#)
- [Section 12.2.2, “Setting Autoprep Facts in the Orchestration Console,” on page 126](#)
- [Section 12.2.3, “Applying Autoprep Facts,” on page 129](#)
- [Section 12.2.4, “Example Autoprep Scenarios,” on page 129](#)
- [Section 12.2.5, “Known Autoprep Limitations,” on page 130](#)

12.2.1 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 hostname 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.

IMPORTANT: If you plan to prepare virtual machines that use LVM as their volume manager on a SUSE Xen or KVM VM host, and if that VM host also uses LVM as its volume manager, you cannot perform autoprep if the VM has an LVM volume with the same name as one already mounted on the VM host. This is because LVM on the VM Host can mount only one volume with the same name.

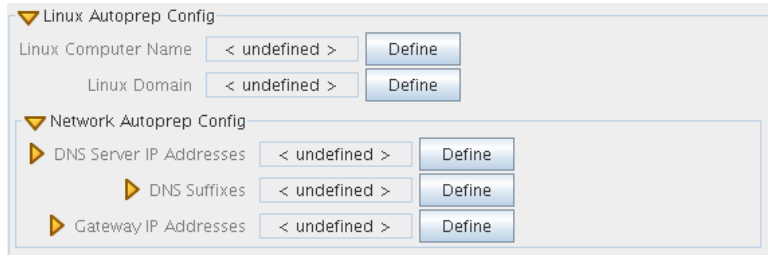
To work around this issue, ensure that the volume names on the VM hosts and virtual machines are different.

12.2.2 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 12-2 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 126](#)
- ♦ [“Network Autoprep Config” on page 127](#)
- ♦ [“Defining Autoprep/Sysprep Network Adapter Facts on the vNIC Object” on page 127](#)

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 make sure 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 box, 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 127](#).

12.2.3 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 *NetIQ 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.

12.2.4 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 autogenerated 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 *NetIQ Cloud Manager Component 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 box, 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 *NetIQ Cloud Manager Procedures Guide*.

12.2.5 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.
- ♦ If a Hyper-V VM is running during the discover process, it fails to discover the `resource.os.family` fact. This prevents the Orchestration Console from displaying the *Sysprep* and *Autoprep* options section on the *Info/Groups* tab for the VM.

NOTE: If the VM not running when the discovery occurs, the hyperv provisioning adapter discovers `resource.os.family` itself.

If you create a template from this VM, the `resource.os.family` fact is discovered and populated on the VM template admin view.

To display the sysprep/autoprep settings on a Hyper-V VM:

1. Shut down the Hyper-V VM that has the problem.
2. From the Explorer Tree, right-click the VM you shut down, then select *Resync State*.
3. In the Orchestration Console, Shift+click the Refresh  icon or restart the Orchestration Console to refresh all of the objects and their facts in the Resource admin view.

13 Using the Cloud Manager Application Server Configuration Tool

After you have installed the NetIQ 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 SUSE Linux Enterprise Server 11 SP2 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 63.

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 chapters provide the detail about the segments of the script:

13.1 Configuring the PostgreSQL Database Connection and Credentials

The NetIQ 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 Postgres to store all of its data.

This section helps you to prepare the information you need to configure the Postgres instance you use for Cloud Manager.

- 1 Make sure you know the information you are prompted to provide during the Postgres configuration:

Information Needed for Configuration	Description
Database server	<p>You need to know the Postgres database server hostname 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.</p> <p>The default is <code>localhost</code>.</p>
Autoconfigure an unconfigured Postgres installation?	<p>If you install a Postgres ORDBMS intended for Cloud Manager but you have not yet configured it, Cloud Manager can autoconfigure it for your environment. Autoconfigure sets up the Postgres authentication method, changes the default postgres user password, and configures the database for local connections only.</p> <p>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.</p> <p>Autoconfigure does the following:</p> <ul style="list-style-type: none"> ◆ Generates a Postgres user password and copies it to <code>/etc/opt/netiq/cloudmanager/etc/pgusr.in</code>. ◆ Creates a database and names it <code>cloudmanager</code>. ◆ Creates a database user and names it <code>cmadmin</code>. ◆ Generates a database password for <code>cmadmin</code> and copies it to <code>/etc/opt/netiq/cloudmanager/etc/com.novell.ncm.backend.connpool.cfg</code> <p>If you choose not to autoconfigure because your Postgres 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.</p> <p>IMPORTANT: Running autoconfigure on a Postgres instance that is already configured causes autoconfigure to fail.</p>
Database server port	<p>You need to know the port that your Postgres server uses for outside communication.</p> <p>The default is <code>5432</code>.</p> <p>NOTE: This documentation does not discuss the configuration for the Postgres server.</p>

Information Needed for Configuration	Description
Create a new Postgres database? (You can choose to use an existing database instead of creating an new one.)	<p>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:</p> <ul style="list-style-type: none"> ♦ 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 <code>cloudmanager</code>. ♦ Database User Name: A user to be created who must have read/write permissions to the database. The default is <code>cmadmin</code>. ♦ Database User Password: The password that the database user (designated above) uses to log in to the database.
Use an existing Postgres database	<p>If you want to use an existing Postgres 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:</p> <ul style="list-style-type: none"> ♦ Database Name: An arbitrary name you specify to identify the database. The default is <code>cloudmanager</code>. ♦ 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 NetIQ 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 Postgres configuration. Use the information you collected in [Step 1 on page 131](#) as the script prompts you.

The configuration tool checks the database server and the database instance you specify, using the newly-defined credentials to make sure that the database instance and the database user can be created.

Following the Postgres configuration, continue with [Chapter 13.3, "Configuring Cloud Manager to Use Authentication Sources," on page 138](#) to configure the authentication sources you want to use with Cloud Manager.

13.2 Configuring the PostgreSQL Database Connection and Credentials

The NetIQ 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 Postgres to store all of its data.

This section helps you to prepare the information you need to configure the Postgres instance you use for Cloud Manager.

- 1 Make sure you know the information you are prompted to provide during the Postgres configuration:

Information Needed for Configuration	Description
Database server	You need to know the Postgres database server hostname 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 <code>localhost</code> .

Information Needed for Configuration	Description
Autoconfigure an unconfigured Postgres installation?	<p>If you install a Postgres ORDBMS intended for Cloud Manager but you have not yet configured it, Cloud Manager can autoconfigure it for your environment. Autoconfigure sets up the Postgres authentication method, changes the default postgres user password, and configures the database for local connections only.</p> <p>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.</p> <p>Autoconfigure does the following:</p> <ul style="list-style-type: none"> ♦ Generates a Postgres user password and copies it to <code>/etc/opt/netiq/cloudmanager/etc/pgusr.in</code>. ♦ Creates a database and names it <code>cloudmanager</code>. ♦ Creates a database user and names it <code>cmadmin</code>. ♦ Generates a database password for <code>cmadmin</code> and copies it to <code>/etc/opt/netiq/cloudmanager/etc/com.novell.ncm.backend.connpool.cfg</code> <p>If you choose not to autoconfigure because your Postgres 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.</p> <p>IMPORTANT: Running autoconfigure on a Postgres instance that is already configured causes autoconfigure to fail.</p>
Database server port	<p>You need to know the port that your Postgres server uses for outside communication.</p> <p>The default is 5432.</p> <p>NOTE: This documentation does not discuss the configuration for the Postgres server.</p>

Information Needed for Configuration	Description
Create a new Postgres database? (You can choose to use an existing database instead of creating an new one.)	<p>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:</p> <ul style="list-style-type: none"> ♦ 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 <code>cloudmanager</code>. ♦ Database User Name: A user to be created who must have read/write permissions to the database. The default is <code>cmadmin</code>. ♦ Database User Password: The password that the database user (designated above) uses to log in to the database.
Use an existing Postgres database	<p>If you want to use an existing Postgres 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:</p> <ul style="list-style-type: none"> ♦ Database Name: An arbitrary name you specify to identify the database. The default is <code>cloudmanager</code>. ♦ 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 NetIQ 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 Postgres configuration. Use the information you collected in [Step 1 on page 131](#) as the script prompts you.

The configuration tool checks the database server and the database instance you specify, using the newly-defined credentials to make sure that the database instance and the database user can be created.

Following the Postgres configuration, continue with [Chapter 13.3, "Configuring Cloud Manager to Use Authentication Sources," on page 138](#) to configure the authentication sources you want to use with Cloud Manager.

13.3 Configuring Cloud Manager to Use Authentication Sources

The instructions in this section assume that you have already used the configuration tool to configure the Postgres database use by the NetIQ Cloud Manager Application Server, as described in [Chapter 13.1, "Configuring the PostgreSQL Database Connection and Credentials," on page 131](#).

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 [Postgres 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 (NetIQ 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 13.3.1, “Configuring Authentication to an LDAP Directory,” on page 139](#)
- ♦ [Section 13.3.2, “Configuring Authentication through an NCSS Director,” on page 141](#)
- ♦ [Section 13.3.3, “Configuring LDAP Plus NCSS Authentication,” on page 143](#)
- ♦ [Section 13.3.4, “Configuring Authentication to Novell Access Manager,” on page 147](#)

13.3.1 Configuring Authentication to an LDAP Directory

The NetIQ Cloud Manager administrator can choose to authenticate users through a supported Lightweight Directory Access Protocol (LDAP) directory service, either Microsoft Active Directory or Novell 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 13.3.2, “Configuring Authentication through an NCSS Director,” on page 141](#), [Section 13.3.3, “Configuring LDAP Plus NCSS Authentication,” on page 143](#), or [Section 13.3.4, “Configuring Authentication to Novell Access Manager,” on page 147](#).

- 1 Make sure 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 Novell eDirectory or Microsoft Active Directory.
LDAP host address	<p>This is the address (DNS name or IP address) of the LDAP host that Cloud Manager can connect to for authentication.</p> <p>If you chose to use SSL with LDAP, this address should match the subject of the certificate issued for the LDAP host.</p> <p>The configuration tool immediately validates this address when you specify it.</p>

Information Needed for LDAP Configuration	Description
LDAP port	<p>Designate the port where you want the LDAP server to listen for communication from Cloud Manager.</p> <p>If you are using SSL, the default port is 636. If you chose not to use SSL, the default port is 389.</p>
Path to SSL certificate on LDAP server	<p>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</p> <p>You need to use this setting only if you want to use SSL with the LDAP authentication.</p>
LDAP read-only user DN	<p>Specify the distinguished name (DN) of an existing LDAP read-only user who has read access to the LDAP directory.</p> <p>This user must have All Attribute read rights to the area of the directory that is to be used for Cloud Manager.</p>
LDAP read-only user's password	<p>Specify the password for the LDAP read-only user.</p> <p>When you specify the user password, the configuration tool immediately attempts an SSL authentication to validate the existence of this user and password.</p>
Cloud Manager LDAP user DN	<p>Specify the DN of an existing LDAP user whom you want to designate as the Cloud Manager administrator.</p> <p>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.</p> <p>Make sure that the <code>mail</code> attribute is set for this user in LDAP.</p>
LDAP DN of NCM Users	<p>Specify the DN of the LDAP container where the users whom you want to log in to Cloud Manager already exist.</p> <p>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.</p> <p>Make sure that all users, regardless of their context in this container, have their email domain configured prior to logging into the Application Console.</p> <p>NOTE: You can use the Cloud Manager Application Console later to import users who do not currently exist in this DN.</p>

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 Postgres database](#), the tool displays the following text:

Authentication Type

- 1) LDAP
- 2) NCSS
- 3) LDAP plus NCSS
- 4) 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 13.4, “Installing and Configuring Other Cloud Manager Feature Settings,”](#) on page 148.

13.3.2 Configuring Authentication through an NCSS Director

The NetIQ Cloud Manager administrator can choose to authenticate users logging in with their email addresses through a supported NetIQ Cloud Security Service (NCSS) server. NCSS should already be installed in your environment.

If you choose to let users authenticate through NCSS, you must also use the Secure Sockets Layer (SSL) protocol with it.

This section helps you to prepare the information you need to configure NCSS for Cloud Manager authentication. If you want to use some other authentication service, see [Section 13.3.1, “Configuring Authentication to an LDAP Directory,”](#) on page 139, [Section 13.3.3, “Configuring LDAP Plus NCSS Authentication,”](#) on page 143, or [Section 13.3.4, “Configuring Authentication to Novell Access Manager,”](#) on page 147.

- 1 Make sure you know the information you’ll be prompted to provide during the NCSS authentication configuration:

Information Needed to Configure Authentication to NCSS Director	Description
DNS Address of the NCSS Director service	Specify the DNS name of the server that hosts the NCSS Director service. This address should match the address on the SSL certificate that was issued for the server.
Path to the SSL Certificate of the NCSS Director server	<p>Specify the path in the file system where the SSL certificate resides. This certificate must be in DER format.</p> <p>If no SSL certificate exists, you can create one by visiting the NCSS Web page in your browser. You can use your browser tools to export the certificate. Remember that it must be in DER format.</p> <p>For more information, see Retrieving the Public Certificate of the LDAP Server (https://www.netiq.com/documentation/cloudsecurityservice/install/data/bqc05g1.html#bqg2k8a) in the <i>NetIQ Cloud Security Service 1.5 Installation Guide</i>.</p>

Information Needed to Configure Authentication to NCSS Director	Description
Cloud Manager Administrator user name	<p>Specify the initial user name that you want to designate as the Cloud Manager administrator.</p> <p>This should be the new administrator's login name or Common Name (CN) and must already exist in your LDAP directory.</p> <p>This value is not validated during the configuration. You must be certain that you specify the value correctly so that users can log in through NCSS.</p>
Cloud Manager Administrator email address	<p>Specify the email address of the user you want to be the Cloud Manager administrator.</p> <p>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.</p> <p>Cloud Manager uses this email address to determine the administrative permissions to apply to the user.</p>

As you continue running the configuration tool (/opt/netiq/cloudmanager/configurator/config) following the [configuration of the Postgres database](#), the tool displays the following text:

Authentication Type

- 1) LDAP
- 2) NCSS
- 3) LDAP plus NCSS
- 4) NAM

Selection:

- 2** Specify 2 (NCSS) as the authentication type you want to configure.
- 3** Follow the prompts and use the information you gathered in [Step 1](#) to complete this segment of the configuration.

After the NCSS authentication configuration, continue with [Chapter 13.4, "Installing and Configuring Other Cloud Manager Feature Settings,"](#) on page 148.

13.3.3 Configuring LDAP Plus NCSS Authentication

The NetIQ Cloud Manager administrator can choose to authenticate tenant customers through a supported NetIQ Cloud Security Service (NCSS) server that redirects the customers back to their own LDAP source for authentication credentials.

With this authentication option, Cloud Manager users of various roles store their credentials in the LDAP directory that you specify during configuration. For example, if you are a cloud service provider, you can set up your own enterprise LDAP structure for logging in to Cloud Manager, but your customers can use their own LDAP structures—preconfigured within NCSS—to authenticate to Cloud Manager.

This section helps you to prepare the information you need to configure LDAP plus NCSS for Cloud Manager authentication.

If you want to use another authentication service, see [Section 13.3.1, “Configuring Authentication to an LDAP Directory,” on page 139](#), [Section 13.3.3, “Configuring LDAP Plus NCSS Authentication,” on page 143](#), or [Section 13.3.4, “Configuring Authentication to Novell Access Manager,” on page 147](#).

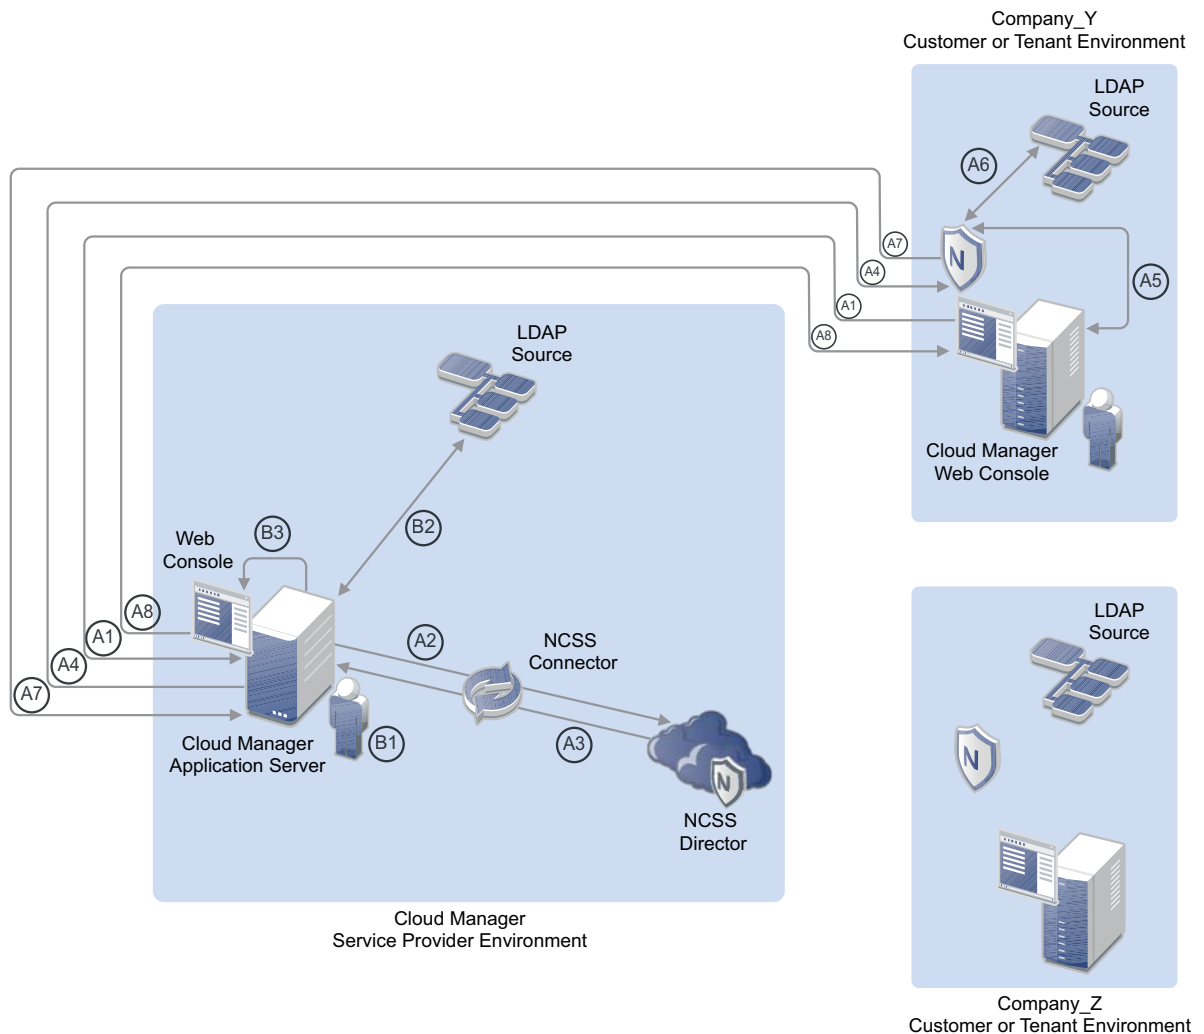
- ♦ [“LDAP Plus NCSS Authentication Concepts” on page 144](#)
- ♦ [“Configuring LDAP plus NCSS Authentication” on page 145](#)

If you want to learn more about NetIQ Cloud Security Service, see the [NetIQ Cloud Security Service 1.5 documentation Web site \(https://www.netiq.com/documentation/cloudsecurityservice/\)](https://www.netiq.com/documentation/cloudsecurityservice/).

LDAP Plus NCSS Authentication Concepts

The following diagram illustrates the process of LDAP plus NCSS authentication from the perspective of the customer or “tenant” on the Cloud Manager Service Provider after you provide the required configuration information.

Figure 13-1 LDAP Plus NCSS Authentication



Customer or “Tenant” Authentication

Stage A1: Customer Steve at Company_Y has bought tenant rights in Jim’s Cloud Manager Service Provider environment. Steve opens a Web browser and loads the URL to the Cloud Manager Application Console. Steve’s Web browser opens the Cloud Manager Web portal, which connects to the Cloud Manager Application Server.

Stages A2, A3: The Application Server uses a preconfigured NCSS connector to communicate to the NCSS Director. The NCSS Director recognizes the login request and redirects that request back to the Cloud Manager Application Server.

Stage A4: The Application Server sends login requirements to NCSS components that were previously installed in the Company_Y environment and directs the Application Console to display open login fields.

Stage A5: NCSS components at Company_Y recognize that an LDAP login is required and establish a link to the Company_Y LDAP source. Steve at Company_Y sees the login field waiting for input on the Cloud Manager Application Console. He enters his username and password.

Stages A6, A7: NCSS components at Company_Y check the customer's LDAP source to validate Steve's login credentials. The NCSS components send the tested credentials back to the Cloud Manager Server, which validates them through the NCSS Director.

Stage A8: The Cloud Manager Server recognizes the privileges that Steve has, based on the context of the email address information listed in his LDAP Source. The Cloud Manager Server opens the relevant Web page for Steve in its Application Console.

Service Provider Administrator Authentication

Stage B1: Service Provider Jim is the Cloud Manager administrator. He wants to use NetIQ Cloud Manager to provide business services to his customers, including Steve at Company_Y, so he opens a Web browser and loads a special URL provided to him by the Cloud Manager Application Server.

Stage B2: Jim's Web browser opens the Cloud Manager Application Console, which is directly connected to the Cloud Manager Server. The Cloud Manager server recognizes the special URL as a Cloud Manager administrator login request, so it establishes a connection to Jim's LDAP source and directs its Application Console to display open login fields for a potential Cloud Manager Administrator.

Stage B3: Jim sees the login field waiting for input on the Cloud Manager Application Console, so he enters his username and password.

The Cloud Manager Server also recognizes the privileges that Jim has, based on the context of the email address information listed in his LDAP Source. The Cloud Manager Server opens the relevant Web page for Jim in its Application Console.

Configuring LDAP plus NCSS Authentication

- 1 Make sure you know the information you'll be prompted to provide during the LDAP plus NCSS authentication configuration:

Information Needed to Configure Authentication	Description
NCSS Director Configuration	
DNS Address of the NCSS Director service	Specify the DNS name of the server that hosts the NCSS Director service. This address should match the address on the SSL certificate that was issued for the server.
Path to the SSL Certificate of the NCSS Director server	Specify the path in the file system where the SSL certificate resides. This certificate must be in DER format. If no SSL certificate exists, you can create one by visiting the NCSS Web page in your browser. You can use your browser tools to export the certificate. Remember that it must be in DER format.
LDAP Configuration	

Information Needed to Configure Authentication	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 Novell eDirectory or Microsoft Active Directory.
LDAP host address	<p>This is the address (DNS name or IP address) of the LDAP host that Cloud Manager can connect to for authentication.</p> <p>If you chose to use SSL with LDAP, this address should match the subject of the certificate issued for the LDAP host.</p> <p>The configuration tool immediately validates this address when you specify it.</p>
LDAP port	<p>Designate the port where you want the LDAP server to listen for communication from Cloud Manager.</p> <p>If you are using SSL, the default port is 636. If you chose not to use SSL, the default port is 389.</p>
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.
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.
LDAP read-only user's password	<p>Specify the password for the LDAP read-only user.</p> <p>When you specify the user password, the configuration tool immediately attempts an SSL authentication to validate the existence of this user and password.</p>
Cloud Manager LDAP user DN	<p>Specify the DN of an existing LDAP user whom you want to designate as the Cloud Manager administrator.</p> <p>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.</p> <p>Make sure that the <code>mail</code> attribute is set for this user in LDAP.</p>
LDAP DN of NCM Users	<p>Specify the DN of the LDAP container where the users whom you want to log in to Cloud Manager already exist.</p> <p>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.</p> <p>Make sure that all users, regardless of their context in this container, have an email domain configured prior to logging into the Application Console.</p> <p>NOTE: You can use the Cloud Manager Application Console later to import users who do not currently exist in this DN.</p>

As you continue running the configuration tool (`/opt/netiq/cloudmanager/configurator/config`) following the [configuration of the Postgres database](#), the tool displays the following text:

Authentication Type

- 1) LDAP
- 2) NCSS
- 3) LDAP plus NCSS
- 4) NAM

Selection:

- 2 Specify 3 (LDAP plus NCSS) as the Authentication Type you want to configure.
- 3 Follow the prompts and use the information you gathered in [Step 1](#) to complete this segment of the configuration.

After the LDAP plus NCSS authentication configuration, continue with [Chapter 13.4, “Installing and Configuring Other Cloud Manager Feature Settings,”](#) on page 148.

13.3.4 Configuring Authentication to Novell Access Manager

The NetIQ Cloud Manager administrator can choose to authenticate customers through Novell Access Manager (NAM).

This section helps you to prepare the information you need to configure Cloud Manager authentication through Novell Access Manager. If you want to use some other authentication service, see [Section 13.3.1, “Configuring Authentication to an LDAP Directory,”](#) on page 139, [Section 13.3.2, “Configuring Authentication through an NCSS Director,”](#) on page 141, or [Section 13.3.3, “Configuring LDAP Plus NCSS Authentication,”](#) on page 143.

If you want to learn more about Novell Access Manager, see the [Novell Access Manager 3.1 SP3 documentation Web site](#) (<http://www.novell.com/documentation/novellaccessmanager313/>).

- 1 Make sure 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	<p>Specify the initial user name that you want to designate as the Cloud Manager administrator.</p> <p>This should be the new administrator’s login name or Common Name (CN) and must already exist in your LDAP directory.</p>
Cloud Manager Administrator email address	<p>Specify the email address of the user you want to be the Cloud Manager administrator.</p> <p>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.</p> <p>Cloud Manager uses this email address to determine the administrative permissions to apply to the user.</p>

As you continue running the configuration tool (`/opt/netiq/cloudmanager/configurator/config`) following the [configuration of the Postgres database](#), the tool displays the following text:

Authentication Type

- 1) LDAP
- 2) NCSS
- 3) LDAP plus NCSS
- 4) NAM

Selection:

- 2 Specify 4 (NAM) as the Authentication Type you want to configure.
- 3 Follow the prompts and use the information you gathered in [Step 1 on page 147](#) to complete this segment of the configuration.

After the Novell Access Manager authentication configuration, continue with [Chapter 13.4, “Installing and Configuring Other Cloud Manager Feature Settings,” on page 148](#).

13.4 Installing and Configuring Other Cloud Manager Feature Settings

When you have completed configuring the NetIQ Cloud Manager Application Server to use [your chosen configuration source](#), you must use the NetIQ Cloud Manager configuration tool to install or configure other Cloud Manager features that help you administer Cloud Manager.

- [Section 13.4.1, “Installing the Cloud Manager Application Console,” on page 148](#)
- [Section 13.4.2, “Configuring the Cloud Manager Web Server \(Jetty\),” on page 148](#)
- [Section 13.4.3, “Configuring the Cloud Manager Web Server to Use SSL,” on page 149](#)
- [Section 13.4.4, “Configuring Cloud Manager SMTP Mail Settings,” on page 150](#)
- [Section 13.4.5, “Configuring Cloud Manager System Shell Login Information,” on page 151](#)

13.4.1 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 [NetIQ Cloud Manager Procedures Guide](#).

13.4.2 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 NetIQ 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 <code>jetty.xml</code> for security purposes. You can re-enable the port by uncommenting the relevant section of the file.

13.4.3 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', NetIQ Cloud Manager will use a self-signed certificate.

Use existing certificate? (y/n):

Make sure 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: <code>/home/jdoe/cloudmgr/newcert.pem</code> SSL is required if you want to use NCSS with Cloud Manager. If no SSL certificate exists, you can create one by using OpenSSL (http://www.novell.com/communities/node/4048/generating-edirectory-server-certificate-using-openssl-tool) or YaST (http://www.novell.com/documentation/sles11/book_security/data/sec_security_yast_ca_module.html). Use your browser tools to export the certificate.

Information Needed to Configure an Imported SSL Certificate	Description
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: <code>/home/jdoe/cloudmgr/newkey.pem</code>
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.

13.4.4 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, make sure 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: <code>smtp.example.test</code> .
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.

13.4.5 Configuring Cloud Manager System Shell Login Information

As the system administrator, you have access to the inner workings of NetIQ 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 doesn't need to correlate to any directory password. It is stored in the <code>users.properties</code> file located in <code>/etc/opt/netiq/cloudmanager/etc</code> .

IV Advanced Installation and Integration Topics

PartIntro

- ♦ [Chapter 14, “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support,” on page 155](#)
- ♦ [Chapter 15, “Installing and Configuring the Orchestration Agent for Xen VM Deployment in a SLES HAE Cluster,” on page 169](#)
- ♦ [Chapter 16, “Configuring the Orchestration Server to Use an Audit Database,” on page 177](#)
- ♦ [Chapter 17, “Integrating the Orchestration Server with a Sentinel Collector,” on page 189](#)
- ♦ [Chapter 18, “Configuring Secure Authentication Sources to Communicate with Cloud Manager,” on page 199](#)

14 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.

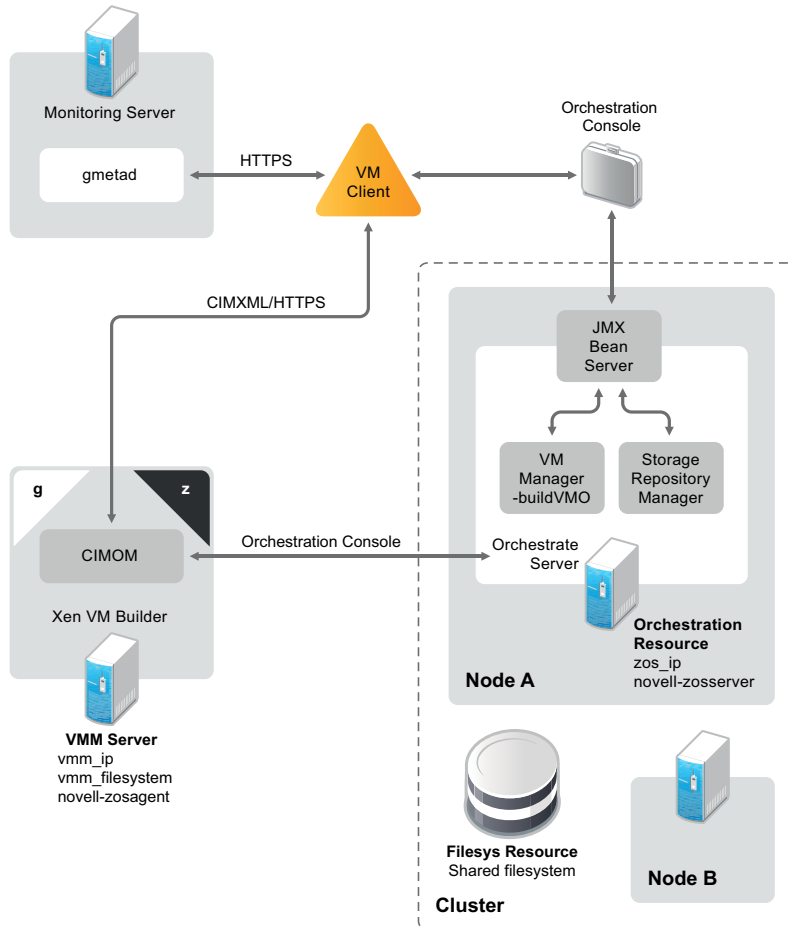
This guide describes how to configure the Cloud Manager Orchestration Server in a high availability SUSE Linux cluster and how to provide both service-level restart for the Orchestration Server and failover among the physical servers of a SUSE Linux cluster to ensure that the server remains available and responsive to the infrastructure that it manages.

- ♦ [Section 14.1, “Overview,” on page 155](#)
- ♦ [Section 14.2, “Orchestration Server Failover Behaviors,” on page 156](#)
- ♦ [Section 14.3, “Installing the Orchestration Server to a SLES 11 Pacemaker Cluster Environment,” on page 157](#)
- ♦ [Section 14.4, “Configuring the Orchestration Server for High Availability,” on page 162](#)
- ♦ [Section 14.5, “Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster,” on page 166](#)
- ♦ [Section 14.6, “Creating the Server Cluster Resource Group,” on page 167](#)
- ♦ [Section 14.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 167](#)
- ♦ [Section 14.8, “Installing and Configuring other Orchestration Components to the High Availability Grid,” on page 167](#)
- ♦ [Section 14.9, “High Availability Best Practices,” on page 168](#)

14.1 Overview

The following figure illustrates how the Orchestration Server is configured for use in a high availability environment.

Figure 14-1 The Orchestration Server in a Clustered, High Availability Environment



14.2 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 14.2.1, “Use Case 1: Orchestration Server Failover,” on page 156](#)
- ♦ [Section 14.2.2, “Use Case 2: VM Builder Behavior at Orchestration Server Failover and Failback,” on page 157](#)
- ♦ [Section 14.2.3, “Use Case 3: Monitoring Behavior at Orchestration Server Failover and Failback,” on page 157](#)

14.2.1 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 14.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 167](#).

When the Orchestration Server fails over to a new node, the Orchestration Agents reauthenticate with the new Orchestration Server instance.

14.2.2 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.

14.2.3 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).

14.3 Installing the Orchestration Server to a SLES 11 Pacemaker Cluster Environment

This section includes information to help you install Orchestration Server components in a high availability SLES 11 SP2 environment. The sequence below is the supported method for configuring this environment.

1. [Section 14.3.1, “Meeting the Prerequisites,” on page 158](#)
2. [Section 14.3.2, “Installing the SLES 11 SP2 High Availability Pattern,” on page 159](#)
3. [Section 14.3.3, “Configuring SLES 11 Nodes with Time Synchronization and Installing Pacemaker to Each Node,” on page 159](#)
4. [Section 14.3.4, “Setting Up OCFS2 on SLES 11 SP2,” on page 161](#)
5. [Section 14.3.5, “Installing the Orchestration Server on the First Clustered SLES 11 Node,” on page 161](#)

You also need to install and configure the Orchestration Agent for the SLES 11 SP2 High Availability Extension (Pacemaker) cluster environment. You can find information to help you do this in [Chapter 15, “Installing and Configuring the Orchestration Agent for Xen VM Deployment in a SLES HAE Cluster,” on page 169](#).

NOTE: Upgrading from earlier versions of Cloud Manager Orchestration to a high availability environment is supported. For more information, see [Chapter 20, “Upgrading Cloud Manager Orchestration Components,” on page 217](#).

14.3.1 Meeting the Prerequisites

The environment where the Orchestration Server is installed must meet the 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 158](#)
- ♦ [“Software Requirements for Creating a High Availability Environment” on page 158](#)

Hardware Requirements for Creating a High Availability Environment

The following hardware components are required for creating a high availability environment for the Orchestration Server:

- ♦ A minimum of two SLES 11 SP2 physical servers, each having dual network interface cards (NICs). These servers are the nodes of the cluster where the Orchestration Server is installed and are a key part of the high availability infrastructure.
- ♦ A Fibre Channel or iSCSI Storage Area Network (SAN) or network storage
- ♦ A STONITH 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”](#) (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_fencing.html) in the *SLES 11 High Availability Guide* (http://www.novell.com/documentation/sle_ha/book_sleha/data/book_sleha.html).

Software Requirements for Creating a High Availability Environment

The following software components are required for creating a high availability environment for the Cloud Manager Orchestration Server:

- ♦ The high availability pattern on the SLES 11 SP2 High Availability Environment (HAE) RPM install source, available for download in a [32-bit version](#) (<http://download.novell.com/Download?buildid=zacLblosaRQ~>) and in a [64-bit version](#) (<http://download.novell.com/Download?buildid=9xvsJDAs04~>).

The SLES 11 source includes Oracle Cluster File System 2 (OCFS2), a parallel cluster file system that offers concurrent access to a shared file system. See [Section 14.3.4, “Setting Up OCFS2 on SLES 11 SP2,” on page 161](#) for more information.

SLES 11 SP2 HAE integrates these open source storage technologies (Pacemaker and OCFS) in a high availability installation pattern, which, when installed and configured, is known as the High Availability Storage Infrastructure. This combined technology automatically shares cluster configuration and coordinates cluster-wide activities to ensure predictable administration of storage resources for shared-disk-based clusters.

- ♦ The Pacemaker software package, which is a high availability resource manager that supports multinode failover. This should include all available online updates installed to all nodes that will be part of the Pacemaker cluster. You can download this cluster resource manager at the [Pacemaker project download site](#) (<http://www.clusterlabs.org/doc/>).
- ♦ DNS is installed on the nodes of the cluster for resolving the cluster hostname to the cluster IP.
- ♦ The Orchestration Server is installed on all nodes of the cluster. A two-node or three-node configuration is recommended.
- ♦ (Optional) The Cloud manager Monitoring Server installed on a non-clustered server.

14.3.2 Installing the SLES 11 SP2 High Availability Pattern

The High Availability Environment ISO install pattern is included in the distribution of the SLES HAE 11 SP2 ISO. Use YaST2 (or the command line, if you prefer) to install the packages that are associated with the high availability pattern to each physical node that is to participate in the Orchestration Server cluster.

NOTE: The high availability pattern is included on the SLES HAE 11 SP2 install source, not the Cloud Manager install source.

The packages associated with high availability include:

- ♦ drbd (Distributed Replicated Block Device)
- ♦ EVMS high availability utilities
- ♦ The Pacemaker subsystem for high availability on SLES
- ♦ The Pacemaker CIM provider
- ♦ A monitoring daemon for maintaining high availability resources that can be used by Pacemaker
- ♦ A plug-in and interface loading library used by Pacemaker
- ♦ An interface for the STONITH device
- ♦ OCFS2 GUI tools
- ♦ OCFS2 Core tools

The packages that must be installed, at a minimum, include:

- ♦ `OCFS2-tools-o2cb`
- ♦ `yast2-cluster`
- ♦ `libglue-devel`
- ♦ `sle-hae-release`

All other dependencies are installed by default.

For more information, see “Installation and Basic Setup with YaST” (http://www.novell.com/documentation/sles11/book_sle_deployment/data/cha_inst.html) in the *SUSE Linux Enterprise High Availability Extension Administration Guide*.

14.3.3 Configuring SLES 11 Nodes with Time Synchronization and Installing Pacemaker to Each Node

When 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 160
- ♦ “Configuring Pacemaker” on page 160

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 SLES HAE 11 SP2. Use the [Time Synchronization with NTP](http://www.novell.com/documentation/sles11/book_sle_admin/data/cha_netz_xntp.html) (http://www.novell.com/documentation/sles11/book_sle_admin/data/cha_netz_xntp.html) instructions in the *SUSE Linux Enterprise Server 11 High Availability Extension Administration Guide* to help you configure each cluster node with NTP.

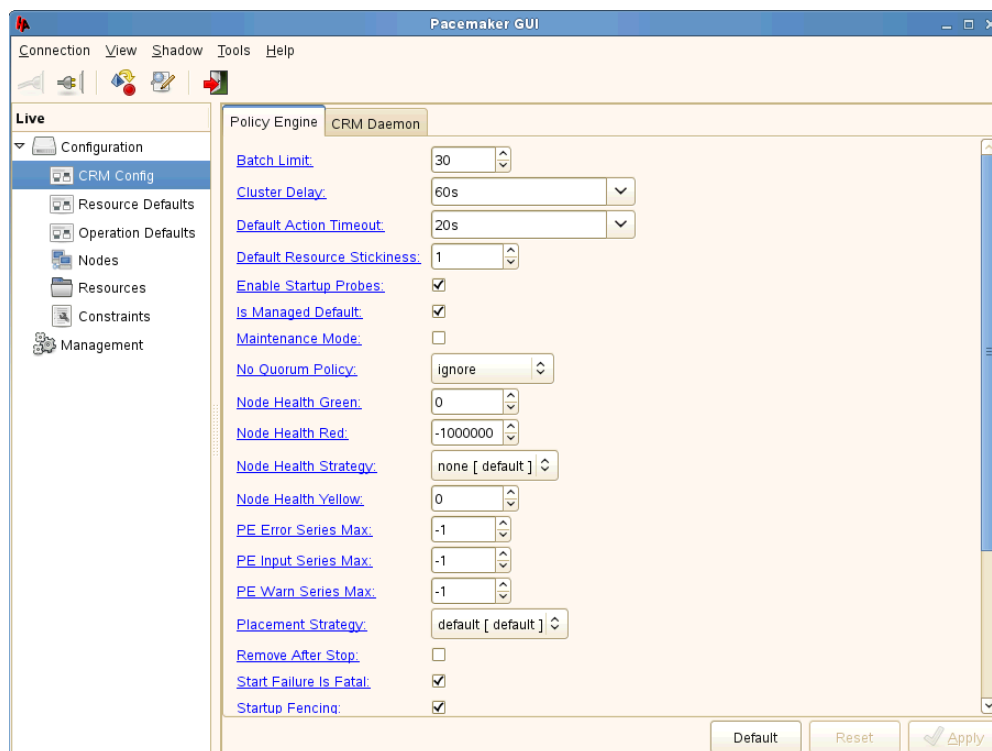
Configuring Pacemaker

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 high availability pattern on the SLES HAE 11 SP2 install source. For detailed information about configuring Pacemaker, see [Configuring and Managing Cluster Resources \(GUI\)](http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_configuration_gui.html) (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_configuration_gui.html) in the *SLES 11 High Availability Guide* (http://www.novell.com/documentation/sle_ha/book_sleha/data/book_sleha.html).

An important value you need to specify in order for Pacemaker to be enabled for high availability is configured in the *Default Action Timeout* field on the settings page of the Pacemaker console.

Figure 14-2 The Main Settings Page in the Pacemaker Graphical Interface



The value in this field controls how long Pacemaker waits for services to start. The default value is 20 seconds. The Orchestration Server requires more time than this to start. We recommend that you specify the value in this field at 120s. More time might be required if your Orchestration Server grid is very large.

14.3.4 Setting Up OCFS2 on SLES 11 SP2

OCFS2 is a general-purpose journaling file system that is fully integrated in the Linux 2.6 and later kernel that ships with SLES 11 SP2. OCFS2 allows you to store application binary files, data files, and databases on devices using network 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.

To set up the high availability environment for the Orchestration Server, you need to first install the High Availability pattern in YaST (this includes the `ocfs2-tools-o2cb` and `ocfs2console` software packages) and configure the Pacemaker cluster management system on each physical machine that participates in the cluster, and then provide network storage with OCFS2 where the Orchestration files can be stored. For information on setting up and configuring OCFS2, see “*Oracle Cluster File System 2*” (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_ocfs2.html) in the *SLES 11 High Availability Guide*.

Shared Storage Requirements for Creating a High Availability Environment

The High Availability Extension available in SLES 11 SP2 supports Fibre Channel or iSCSI storage area networks (SANs).

SAN configuration is beyond the scope of this document. For information about setting up a SAN, see the *Oracle Cluster File System 2* (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_ocfs2.html) documentation in the *SLES 11 High Availability Guide*.

IMPORTANT: The Cloud Manager Orchestration Server requires a specific mount point for file storage on the SAN. Use `/zos` for this mount point.


14.3.5 Installing the Orchestration Server on the First Clustered SLES 11 Node

NOTE: As you prepare to install the Cloud Manager Orchestration Server and use it in a high availability environment, make sure that the requirements to do so are met. For more information, see [Chapter 2, “Cloud Manager System Requirements,” on page 19](#).

To install the Orchestration Server packages on the first node of the cluster:

- 1 Log in to the target SLES 11 server as `root`, then open YaST2.
- 2 Download the appropriate NetIQ Cloud Manager ISO to the SLES server.
or
Load the NetIQ Cloud Manager DVD on the SLES server.
- 3 Define the NetIQ 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 box.



- 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 box.
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 165](#).

14.4 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, make sure you reconfigure all of the components that are installed on the system (this is the default).

- ♦ [Section 14.4.1, “Some Considerations When Configuring with the GUI Wizard,” on page 163](#)
- ♦ [Section 14.4.2, “The Configuration Procedure,” on page 164](#)
- ♦ [Section 14.4.3, “Checking the Configuration,” on page 165](#)
- ♦ [Section 14.4.4, “Running the High Availability Configuration Script,” on page 166](#)

When you have configured the SLES 10x or the SLES 11 SP1 Orchestration Server, you need to complete the other items necessary for a high availability setup in the following order:

1. [Section 14.5, “Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster,” on page 166.](#)
2. [Section 14.6, “Creating the Server Cluster Resource Group,” on page 167.](#)
3. [Section 14.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 167](#)
4. [Section 14.8, “Installing and Configuring other Orchestration Components to the High Availability Grid,” on page 167.](#)

14.4.1 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.

14.4.2 The Configuration Procedure

To configure the Orchestration Server for use in a high-availability environment,

- 1 Make sure you are logged in as `root` to run the configuration.
- 2 Make sure you are ready with the information that you'll be prompted for during the configuration procedure (GUI or text-based):

Server Configuration Requirement	Explanation and Action
Configuration Type	<p>Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation.</p> <p>This section discusses standard installation, so specify <code>h</code> (for <code>ha</code> which means "high availability").</p>
Cluster Hostname or IP Address	<p>Specify the fully qualified cluster hostname or the IP address that is used for configuring the Orchestration Server instance in a high availability cluster.</p> <p>The configuration script binds the IP address of the cluster to this server.</p>
Grid Name	<p>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.</p> <p>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.</p>
Administrator User	<p>Specify a name for the Orchestration Server Administrator user.</p> <p>This name is used to log in as the administrator of the Orchestration Server and the objects it manages.</p>
Administrator Password	<p>Specify a password for the Orchestration Administrator user, then retype the password to validate it.</p> <p>You should remember this username for future logins.</p>
Path to License File	<p>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.</p>
Auditing Database	<p>We recommend that you do not install the audit database on this server.</p>
Orchestration Agent Port ¹	<p>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.</p> <p>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.</p> <p>This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the <i>Configure Advanced Settings</i> 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.</p>

Server Configuration Requirement	Explanation and Action
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. <ul style="list-style-type: none"> ♦ Default = <code>yes</code> (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 <code>no</code>, 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. <ul style="list-style-type: none"> ♦ Default = <code>/etc/ssl/servercerts/servercert.pem</code> ♦ Specify the path to the existing TLS certificate.
TLS Server Key ²	Specify the full path to the TLS server private key. <ul style="list-style-type: none"> ♦ Default = <code>/etc/ssl/servercerts/serverkey.pem</code> ♦ 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 165](#).

14.4.3 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 make sure 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: Make sure you configure DNS to resolve the cluster hostname 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 166.](#)

14.4.4 Running the High Availability Configuration Script

Before you run the high availability configuration script, make sure that you have installed the Orchestration Server to a single node of your high availability cluster. For more information, see [Section 14.3.5, “Installing the Orchestration Server on the First Clustered SLES 11 Node,” on page 161](#)

IMPORTANT: The high availability configuration script asks for the mount point on the Fibre Channel SAN. Make sure 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. Make sure that you follow the prompts in the script exactly; do not misidentify a secondary node in the cluster as the primary node.

14.5 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):

14.6 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 failbacks.

When you have installed and configured the nodes in the SLES 11 SP2 cluster and created a cluster resource group, use the Pacemaker tools to start the cluster resource group. For more information, see “[Cluster Management Tools \(http://www.novell.com/documentation/beta/sle_ha/book_sleha/data/cha_ha_management.html\)](http://www.novell.com/documentation/beta/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 14.7, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 167](#)).

14.7 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.

14.8 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 19](#).
- ♦ 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 43](#).

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 hostname 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 43](#).

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).

14.9 High Availability Best Practices

This section includes information that might be useful to users of the NetIQ Cloud Manager Orchestration Server in high availability environments. We anticipate that the contents of the section will expand as the product is adopted and deployed. We encourage your contributions to this document. All comments are tested and approved by product engineers before they appear in this section. Please use the User Comments feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html (<http://www.novell.com/documentation/feedback.html>) and enter your comments there.

- ♦ [Section 14.9.1, “Jobs Using scheduleSweep\(\) Might Need a Start Constraint,” on page 168](#)

14.9.1 Jobs Using scheduleSweep() Might Need a Start Constraint

If you write a custom job that uses the `scheduleSweep()` JDL function to schedule joblets and that are either 1) marked as restartable in a high availability failover situation or 2) scheduled through the Job Scheduler to run at server startup, the job might fail to schedule any joblets and is easily noticeable with a 0 second run time. This is because `scheduleSweep()`, by default, creates joblets only for online nodes.

If the Job runs during failover, resources might not be readily available, so the job ends immediately.

To keep the Job from running until a resource is online, you can use a start constraint. For example, you could add the following to the job policy:

```
<constraint type="start" >
  <gt fact="jobinstance.matchingresources" value="0" />
</constraint>
```

If you implement this constraint, the Job is queued (not started) until at least one resource matches the policy resource constraint.

As alternatives to using the constraint approach, you can:

- ♦ Code in a waiting interval for the required Agents in your Job
- ♦ Using the `schedule()` API for creating Joblets instead of the `scheduleSweep()` function.
- ♦ Choose an alternative set of resources to consider for the `scheduleSweep()`. For more information, see the “ScheduleSpec” API for more details.

15 Installing and Configuring the Orchestration Agent for Xen VM Deployment in a SLES HAE Cluster

The Cloud Manager Orchestration Server provides data center administrators the capability of providing high availability for Xen-based VMs that is comparable to the service level offerings of other hypervisors such as Microsoft Hyper-V and VMware vSphere. This is made possible with the additional installation of SUSE Linux Enterprise Server 11 Support Pack 2 (SLES 11 SP2) High Availability Extension (HAE) cluster stack. Now, the administrator can use the Orchestration Server to discover SUSE Linux high availability clusters, delegate VM provisioning to the cluster, and perform life cycle operations through the cluster manager.


- ♦ [Section 15.1, “Xen Cluster Architecture,” on page 169](#)
- ♦ [Section 15.2, “Installing the Orchestration Agent in a SLES 11 SP1 HAE Xen Cluster,” on page 170](#)
- ♦ [Section 15.3, “Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster,” on page 171](#)
- ♦ [Section 15.4, “Sample Orchestration Agent CIB XML,” on page 174](#)

15.1 Xen Cluster Architecture

The following diagram shows how Cloud Manager 2.x interacts with a SLES HAE cluster. To manage VMs in the cluster, the Cloud Manager Orchestration Server needs to communicate with the cluster stack. This communication happens through an Orchestration Agent, which the administrator configures as a cluster resource using a special configuration script.

When configured, the SLES HAE cluster chooses which cluster node the Agent runs on, just as it does for any other cluster resource. To make sure that the agent has a consistent IP address, the configuration script sets up an IP address resource in a cluster resource group, along with the cluster resource for the Orchestration Agent. If the Orchestration Agent fails over to another cluster node, its cluster IP address moves with it.

Figure 15-1 Cloud Manager Interaction with a SLES HAE Cluster



Assuming that a SLES HAE cluster is correctly installed and configured, setting up the additional cluster resources is relatively uncomplicated. The Orchestration Server administrator installs and configures the Orchestration Agent on each node in the cluster. In a standard agent configuration, this launches the agent, but because this configuration is for a high availability environment, the agent is not started.

To configure the cluster resources for the agent, the administrator runs a special configuration script on a single node in the cluster. This script needs to be run only once. It creates the cluster resource group, a cluster resource for an IP address, and a cluster resource for the Orchestration Agent. The script then starts the cluster resource group in the cluster. Following this configuration, the Orchestration Agent runs in the cluster as a cluster resource where it can be used by the Orchestration Server to communicate with the cluster stack to facilitate VM management in the cluster.

15.2 Installing the Orchestration Agent in a SLES 11 SP1 HAE Xen Cluster

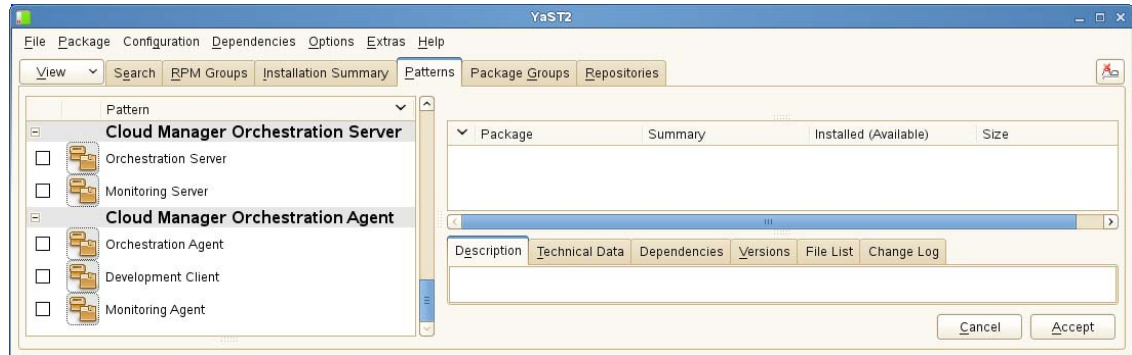
The Orchestration Agent installation pattern can be installed by using the SLES 11 Cloud Manager 2.x installation ISO. Before you install the agent, make sure that the SLES 11 SP2 HAE server machine meets the [prerequisites for agent installation](#). Use the following steps to install a single Orchestration Agent on a single node in a SLES 11 SP2 HAE cluster. You need to repeat these steps on *every* node in the cluster.

- 1 Download the appropriate NetIQ Cloud Manager ISO to the SLES server.

or

Load the NetIQ Cloud Manager DVD on the SLES server.

- 2 Define the NetIQ Cloud Manager ISO or DVD as an add-on product:
 - 2a In the YaST Control Center, click *Software*, then click *Add-On Products*.
 - 2b Click *Add*, select *Local ISO Image* or *DVD*, then follow the prompts to add the product.
- 3 Read and accept the license agreement, then click *Next* to display the Software Selection and System Tasks dialog box.



- 4 Select the Orchestration Agent installation pattern for installation.
- 5 Click *OK* to install the packages.
- 6 When package installation is complete, click *OK* to close the Installed Add-On Products dialog box.
- 7 Install the agent on each node of the SLES 11 SP2 HAE cluster, repeating [Step 1](#) through [Step 6](#) above.

When you complete the pattern installation, use the information in [Section 15.3, “Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster,”](#) on page 171 to configure the Orchestration Agent.

15.3 Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster

After you have installed the Orchestration Agent on each node of the cluster, you need to configure those installations.

- [Section 15.3.1, “Configuring the Agent for the Cluster,”](#) on page 171
- [Section 15.3.2, “Creating the Agent Cluster Resource Group,”](#) on page 173
- [Section 15.3.3, “Removing the Orchestration Agent from a Clustered VM Host,”](#) on page 174

15.3.1 Configuring the Agent for the Cluster

- 1 Make sure you are ready with the information that you are prompted for during the Orchestration Agent 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 agent installation or in an HAE cluster, so specify <i>h</i> (for high availability).

Server Configuration Requirement	Explanation and Action
Cluster Hostname or IP Address	<p>Enter the IP address for use by the Orchestration Agent as it communicates with the Orchestration Server. As a result of the complete configuration process, this IP address is configured as a cluster resource in a cluster resource group and as a resource for the Orchestration Agent.</p> <p>You can use a hostname value instead of an IP address, provided that hostname resolution is set up correctly in your environment.</p> <p>The value that you specify here is used as the default value for two other configuration items: Agent Name and Cluster Bind Address.</p>
Agent Name	<p>This is the name the Orchestration Agent uses to present itself to the Orchestration Server. The agent is represented with this name in the Orchestration Console Explorer tree.</p> <p>The default for this field is the value you specified for <i>Cluster hostname or IP address</i>, but you can modify it to anything you choose (not recommended).</p>
Orchestration Server Hostname or IP Address	This value is required.
Always Implement the Orchestration Server Certificate and Key?	<p>The Agent relies on the Orchestration Server's TLS certificate as verification that it is communicating with the correct Orchestration Server.</p> <p>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.</p>
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.
Cluster Bind Address	<p>This is the IP address the agent should use when connecting to the Orchestration Server. A default value might be derived from the supplied <i>Cluster hostname or IP address</i> as follows:</p> <ul style="list-style-type: none"> ♦ If an IP address was specified, the value is used unchanged. ♦ If a valid DNS hostname was specified (that is, the DNS name is resolveable to an IP address), the IP address associated with the DNS name is set as the default. ♦ If an invalid DNS hostname was specified (that is, the DNS name is unresolveable to an IP address), the field remains empty and you need to enter the IP address manually.
Path to Server Certificate	<p>Specify the path to the Orchestration Server certificate file. The default path is <code>/root/zos_server_cert.pem</code>.</p> <p>NOTE: This configuration parameter is considered an advanced setting for the Orchestration Agent in the GUI Configuration Wizard, but only if you set <i>Provide Existing Orchestration Server Certificate</i> to <i>yes</i>.</p>

2 On any node in the cluster, run the configuration utility:

```

/opt/novell/zenworks/orch/bin/config
or
/opt/novell/zenworks/orch/bin/guiconfig

```

- 3 Follow the prompts in the configuration utility.
- 4 When you have specified all of the information that the configuration tool needs, press Enter to begin the configuration.
- 5 Run the agent configuration utility on each node of the SLES 11 SP2 HAE cluster, repeating [Step 1](#) through [Step 4](#) above.

When you have completed all configurations on all cluster nodes, run the resource group configuration script on exactly one node of the cluster. The information in [Section 15.3.2, “Creating the Agent Cluster Resource Group,” on page 173](#) provides the detail on what the script does and how to run it.

Preventing Corruption in Clustered Storage Repositories

Orchestrate has a well-known VM fact called `resource.vm.reprovisionOnAgentOfflineEvent` that defaults to `False` in order to prevent virtual machines from being reprovisioned when the agent is taken offline on the VM host. This fact can be set to `True`, but it must be understood that if the hosts use a shared cluster storage (such as an OCFS2 repository), setting this to `True` and then taking a VM host's agent offline can cause the VM to become corrupt if the VM is not shut down on the first host it is running on.

In a two-host environment, shutting the agent down on Host1 while leaving the VMs running is the VM running on Host1 would continue to run while it is started on Host2, and both hosts would access the virtual disk files simultaneously without awareness of the other. Both would write data to the shared storage independently, and this would corrupt the virtual disk.

It is strongly recommended that if clustered storage repositories are used, this VM fact is left to the default setting of `False` in order to prevent this type of corruption.

15.3.2 Creating the Agent Cluster Resource Group

The resource group creation script, `zos_agent_ha_resource_group.sh` is located in `/opt/novell/zenworks/orch/bin/ha` with the other configuration tools. Run this script on only one node in the cluster to set up the agent cluster resource group.

The script

- ♦ Creates a CIB definition for the clustered Orchestration Agent.
- ♦ Creates a CIB definition for the agent's clustered IP address.
- ♦ Creates a cluster resource group for the Orchestration Agent and the agent's clustered IP address.
- ♦ Configures the resource stickiness to avoid unnecessary failbacks.

When you run the resource group script, it asks for three parameters:

- ♦ **IP ADDRESS of the cluster hostname used to configure the Orchestration Agent:** Specify the value that you provided for the [Cluster Hostname or IP Address](#) when you configured the Orchestrate Agent on the cluster nodes. This must be an IP address, not a hostname.
- ♦ **NETMASK to be used with the specified IP ADDRESS:** Specify the correct netmask for the cluster IP address you provided for the [Cluster Hostname or IP Address](#) when you configured the Orchestrate Agent on the cluster nodes.
- ♦ **INTERFACE to use when the IP ADDRESS is brought online (optional):** Specify the network interface you want the cluster IP address to bind to.

The collected information is used to create a Cluster Information Base (CIB) XML template for configuring the Orchestration Agent and a cluster IP address in a resource group in the SLES 11 SP2 HAE cluster. The template is called `cluster_zos_agent.xml` and is located in the `/opt/novell/zenworks/orch/bin/ha/` directory. See [Section 15.4, “Sample Orchestration Agent CIB XML,”](#) on [page 174](#) for a sample of this template.

The resource script runs the following command to create the resource group:

```
/usr/sbin/cibadmin -o resources -C -x $XMLFILE
```

The cluster resource group is then brought online by the resource script. For more information about SLES HAE cluster tools, see “[Configuring and Managing Cluster Resources \(GUI\)](http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_configuration_gui.html) (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_configuration_gui.html)” or “[Configuring and Managing Cluster Resources \(Command Line\)](http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_manual_config.html) (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_manual_config.html)” in the *SUSE Linux Enterprise High Availability Extension Administration Guide*.

15.3.3 Removing the Orchestration Agent from a Clustered VM Host

When removing the agent from a clustered VM host, stopping the agent with managed VM resources deployed will cause the resources to fail and the cluster node to be fenced.

To prevent this from happening, use the Pacemaker GUI (`crm_gui`), the HA Web Konsole (Hawk), or the HA command-line tool (`crm`) to make the resource unmanaged. For information on how to accomplish this, see the [SUSE Linux Enterprise High Availability Extension documentation](https://www.suse.com/documentation/sle_ha/book_sleha/data/book_sleha.html) (https://www.suse.com/documentation/sle_ha/book_sleha/data/book_sleha.html).

15.4 Sample Orchestration Agent CIB XML

The following XML is from the file `/opt/novell/zenworks/orch/bin/ha/cluster_zos_agent.xml`. It can be used as an example of how to configure the Orchestration Agent within a SLES 11 SP2 HAE cluster. You must replace the `$CONFIG_ZOS_AGENT_CLUSTER_IP` string with a valid cluster IP address (this is what the `zos_agent_ha_resource_group.sh` script does).

The `resource-stickiness` setting of `+INFINITY` causes the Orchestration Agent to prefer the cluster node where it is currently running, unless a failover occurs (the agent does not continually migrate as the cluster attempts to balance the load).

For example, consider the following scenario:

- An instance of the Orchestration Agent is running on `host1`
- A VM (designated “`vm1`”) is running on `host1`
- You provision another VM (designated as “`vm2`”) to `host1`.

The following events occur in this scenario:

1. The HAE cluster moves the Orchestration Agent instance from `host1` to `host2`.
2. Moving the agent causes the provision job to cancel/fail and the agent to disconnect/log in again to the server.

In this same scenario, the reverse happens if the *Shutdown* action is run on `vm2`:

1. An Orchestration Server job kills the agent instance on `host2`.
2. The HAE cluster moves the agent instance from `host2` to `host1`.

The resource-stickiness setting solves this issue. However, it would also be sufficient to add a resource location constraint on the Orchestration Agent, which would cause it to prefer a desired node in the cluster with a score of +INFINITY.

XML Sample with Resource-Stickiness Setting

```
<group id="novell-zosagent-group">
  <primitive class="ocf" id="novell-zosagent-ip" provider="heartbeat"
type="IPAddr">
    <operations id="novell-zosagent-ip-operations">
      <op id="novell-zosagent-ip-op-monitor-5s" interval="5s"
name="monitor" timeout="20s"/>
    </operations>
    <instance_attributes id="novell-zosagent-ip-instance_attributes">
      <nvpair id="novell-zosagent-instance_attributes-ip" name="ip"
value="151.155.169.78"/>
      <nvpair id="novell-zosagent-instance_attributes-cidr_netmask"
name="cidr_netmask" value="255.255.252.0"/>
      <nvpair id="novell-zosagent-instance_attributes-nic" name="nic"
value="br0"/>
    </instance_attributes>
  </primitive>
  <primitive class="lsb" id="novell-zosagent" type="novell-zosagent">
    <meta_attributes id="novell-zosagent-meta_attributes">
      <nvpair id="novell-zosagent-meta_attributes-resource-stickiness"
name="resource-stickiness" value="+INFINITY"/>
    </meta_attributes>
    <operations id="novell-zosagent-operations">
      <op id="novell-zosagent-op-monitor-15" interval="15" name="monitor"
start-delay="15" timeout="15"/>
    </operations>
  </primitive>
  <meta_attributes id="zos-agent-group-meta_attributes">
    <nvpair id="zos-agent-group-meta_attributes-target-role"
name="target-role" value="started"/>
  </meta_attributes>
</group>
```

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 the Orchestration Server is supported by default on SLES 11 SP2, you can use a PostgreSQL database that ships with SLES 11 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 177](#)
- ♦ [Section 16.2, “Configuring PostgreSQL to Accept Remote Database Connections,” on page 179](#)
- ♦ [Section 16.3, “Logging in Locally to the PostgreSQL Database,” on page 180](#)
- ♦ [Section 16.4, “Creating an Orchestration Server User for the PostgreSQL Database,” on page 180](#)
- ♦ [Section 16.5, “Configuring the Orchestration Server Audit Database on a Separate Host,” on page 180](#)
- ♦ [Section 16.6, “Installing and Configuring the Orchestration Server for Use with a Local PostgreSQL Audit Database,” on page 182](#)
- ♦ [Section 16.7, “Configuring the Audit Database after the Cloud Manager Orchestration Server Is Configured,” on page 185](#)
- ♦ [Section 16.8, “Configuring the Remote Audit Database after the Cloud Manager Orchestration Server Is Configured,” on page 186](#)
- ♦ [Section 16.9, “Modifying Audit Database Tables to Accommodate Long Names,” on page 187](#)
- ♦ [Section 16.10, “Understanding Grid ID Usage in the Audit Database,” on page 187](#)

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 SLES 10 SP2 that is different from the server 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 182.

If the SLES 11 SP2 machine) does not have PostgreSQL packages installed and running, 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 178](#).

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, make sure 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 "164.99.15.64", 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 username 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 username, then press Enter.

This is the same username that was created when PostgreSQL was installed. In most instances, the username 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 by default on SLES 10 SP3, SLES 10 SP4, or SLES 11 SP1, you can use a PostgreSQL database 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 182](#)
- [Section 16.6.2, “Configuring PostgreSQL to Accept Local Database Connections,” on page 183](#)
- [Section 16.6.3, “Logging in Locally to the PostgreSQL Database,” on page 183](#)
- [Section 16.6.4, “Installing and Configuring the Local Orchestration Server Audit Database,” on page 183](#)

16.6.1 Installing the PostgreSQL Package and Dependencies

NOTE: We recommend that you install the PostgreSQL package on a SLES 10 SP3, SLES 10 SP4, or a SLES 11 SP1 server that is different from the server where you install the Cloud Manager Orchestration Server. This ensures 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 179](#).

If your SLES 10 SP3, SLES 10 SP4, or SLES 11 SP1 machine does not have the PostgreSQL package installed and running, 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 install the prepared tables. For more detail about what you will see when you run these commands, see [“Detail” on page 178](#).

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 username 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 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=#
```

NOTE: This is the message you would see if you are logging in to PostgreSQL on a SLES 10 SP3 or SLES 10 SP4 server. If logging in to Postgres on a SLES 11 SP1 server, you would see a message indicating a login to psql 8.3.9.

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 username, then press **Enter**.

This is the same name that was specified when PostgreSQL was installed. In most instances, the username 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.

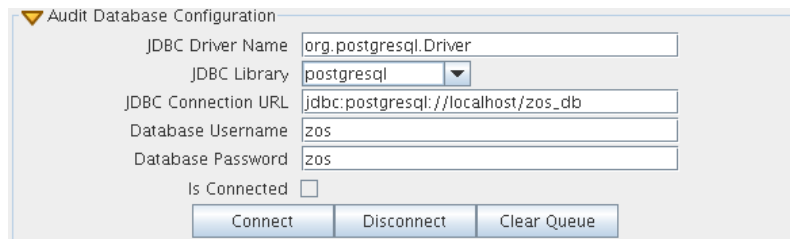
16.7 Configuring the Audit Database after the Cloud Manager Orchestration Server Is Configured

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 username 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 180. In this example, the username is `zos` and the password is `zos`.



- 5 Confirm that the database username and password match the values you replaced in the variables of the `.sql` script. In this example, the username 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.

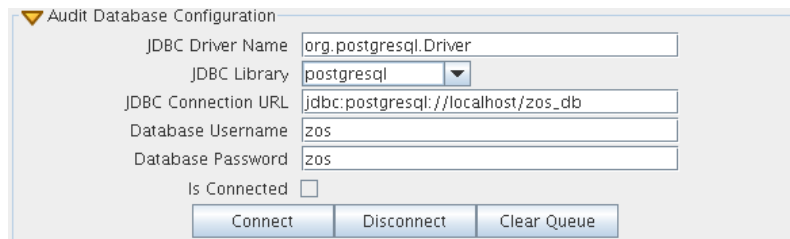
16.8 Configuring the Remote Audit Database after the Cloud Manager Orchestration Server Is Configured

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 username 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 180](#). In this example, the username is `zos` and the password is `zos`.



- 5 Confirm that the database username and password match the values you replaced in the variables of the `.sql` script. In this example, the username 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 `username` 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 Novell Sentinel collector for the Cloud Manager Orchestration Server with your NetIQ Cloud Manager and Novell 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 Novell 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 177.

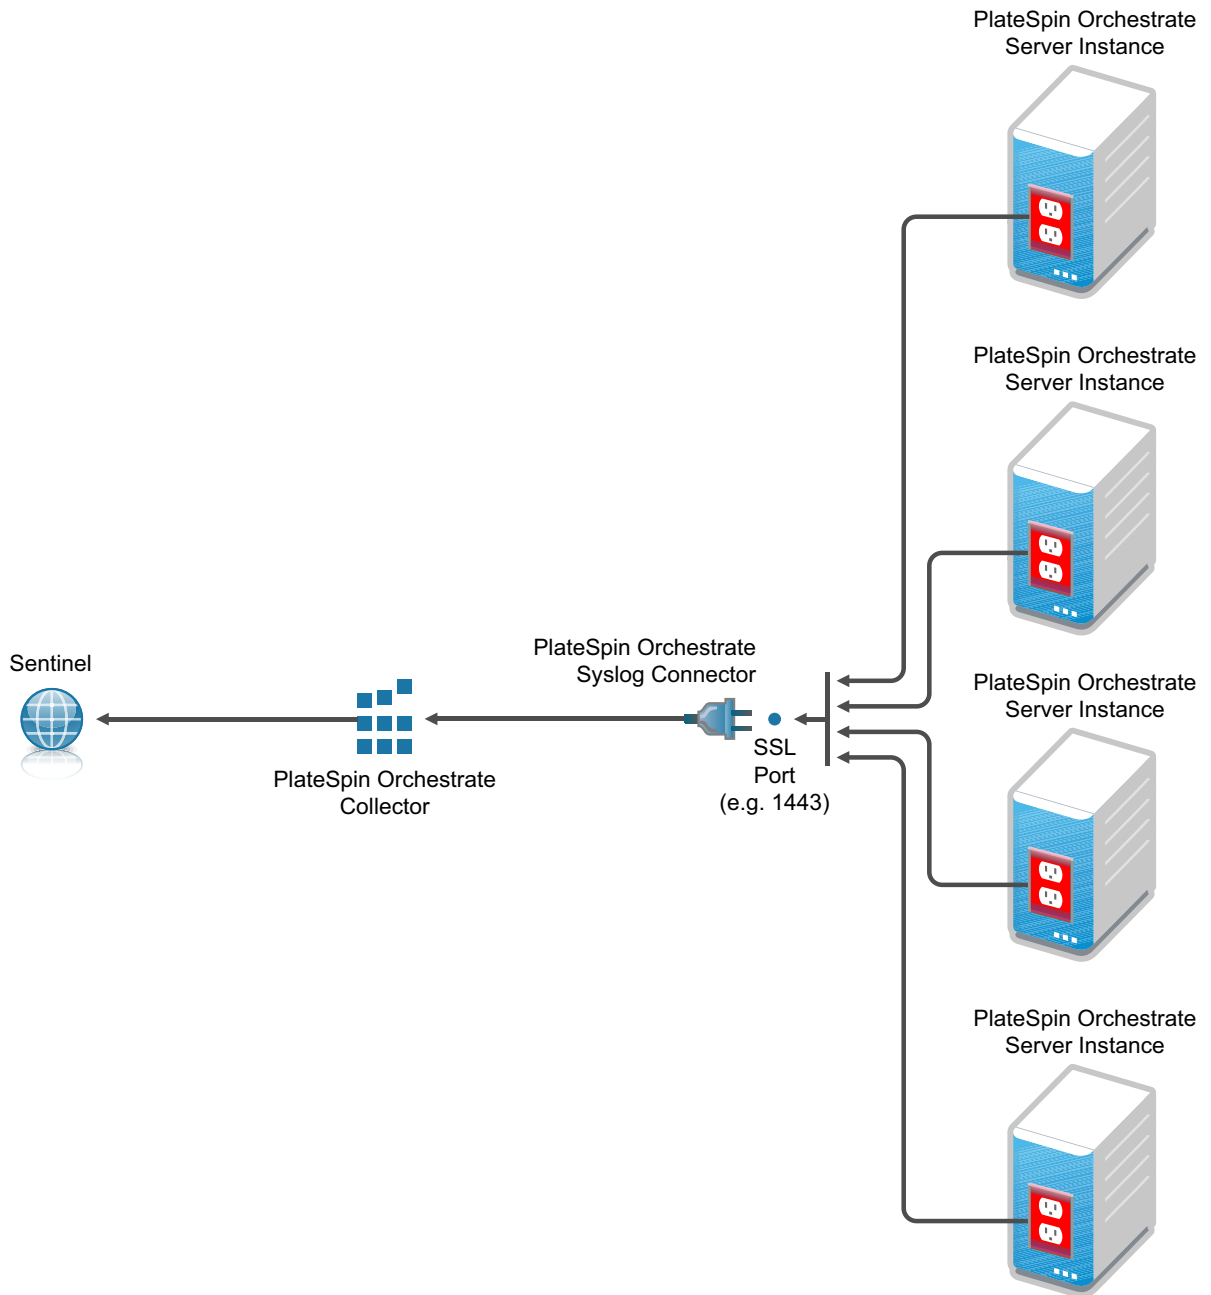
- ♦ [Section 17.1, “Integration Architecture,”](#) on page 189
- ♦ [Section 17.2, “System Requirements,”](#) on page 190
- ♦ [Section 17.3, “Importing and Deploying the Orchestration Server Sentinel Collector Plug-in,”](#) on page 191
- ♦ [Section 17.4, “Connecting the Orchestration Server to the Sentinel Collector Plug-In,”](#) on page 193
- ♦ [Section 17.5, “Verifying the Sentinel Configuration After Connecting to the Orchestration Server,”](#) on page 193
- ♦ [Section 17.6, “Event Classification and Taxonomy Keys,”](#) on page 194
- ♦ [Section 17.7, “Plain Text Visibility of Sensitive Information,”](#) on page 197

17.1 Integration Architecture

Novell 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 Novell Sentinel, see the [Novell Sentinel 6.1 product documentation Web site \(http://www.novell.com/documentation/sentinel61/index.html\)](http://www.novell.com/documentation/sentinel61/index.html).

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 3.0 installed and running on a supported SUSE Linux Enterprise Server.

- ♦ Sentinel 6.1.1.1 (and higher) installed and running.

NOTE: Sentinel 6.1x must be separately purchased from Novell.

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 20.

For more information about the server requirements for Sentinel, see [System Requirements \(http://www.novell.com/documentation/sentinel61/s61_install/data/bgmq7g2.html\)](http://www.novell.com/documentation/sentinel61/s61_install/data/bgmq7g2.html) in the *Novell Sentinel 6.1 Installation Guide*.

- ♦ The appropriate Sentinel Syslog Connector, [available for download from Novell Support \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).
- ♦ The appropriate Orchestration Server/Sentinel Collector Plug-in, available for download on the [Sentinel Collectors download page \(http://support.novell.com/products/sentinel/secure/sentinelplugins.html\)](http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

NOTE: Collector-specific information is available in .pdf format with the download above. Users should consult with Novell 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 Novell 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, make sure 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 Novell 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 [NetIQ Cloud Manager Component 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, make sure 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 the Orchestration Server/Sentinel Collector Plug-in, use the following steps to configure the connection between the plug-in and the 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 hostname 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 *NetIQ Cloud Manager Component 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 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 <code>eventclass</code> 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

Value of the eventclass Taxonomy Key	Events in This Classification	Other Keys Typically Used with This Class of Events
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
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

Value of the eventclass Taxonomy Key	Events in This Classification	Other Keys Typically Used with This Class of Events
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
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 Cloud Security Service (NCSS) and Novell Access Manager (NAM) as identity service tools that you can leverage to let your NetIQ Cloud Manager users securely log in to Cloud Manager.

- ♦ [Section 18.1, “Configuring Cloud Security Service Connectors for Cloud Manager,” on page 199](#)
- ♦ [Section 18.2, “Configuring Novell Access Manager to Work with Cloud Manager,” on page 204](#)

18.1 Configuring Cloud Security Service Connectors for Cloud Manager

NetIQ Cloud Security Service is a Web-based identity and access management gateway that acts as a trusted identity broker between your NetIQ Cloud Manager installation being used as a cloud service provider and the various LDAP infrastructures you might have in place to support and manage your customer's authentication.

Through NCSS Connectors that you configure, your customers can seamlessly become users of Cloud Manager and are enabled to use their accustomed credentials to log in to Cloud Manager with limited access to services specific to their organizations. NCSS and Cloud Manager work in unison to provide this multi-tenant environment where information is separated, secured, and protected. There are four NCSS connectors that you need to configure:

- ♦ [Section 18.1.1, “Creating and Configuring an NCSS Connector for the Cloud Manager Application Console,” on page 200](#)
- ♦ [Section 18.1.2, “Creating and Configuring an NCSS Connector for the Cloud Manager API,” on page 203](#)
- ♦ [Section 18.1.3, “Creating and Configuring an NCSS Connector for the Cloud Manager Tasks API,” on page 203](#)
- ♦ [Section 18.1.4, “Creating and Configuring an NCSS Connector for the Cloud Manager App Services API,” on page 203](#)

The content in this section is not intended as a comprehensive guide to NCSS. You should have already installed NCSS (<https://www.netiq.com/documentation/cloudsecurityservice/install/data/bq5t8g8.html>) and you should be familiar with its capabilities so that you understand the context of this section. For more complete information about NetIQ Cloud Security Service, see the [Cloud Security Service documentation Web site \(https://www.netiq.com/documentation/cloudsecurityservice/\)](https://www.netiq.com/documentation/cloudsecurityservice/).

18.1.1 Creating and Configuring an NCSS Connector for the Cloud Manager Application Console

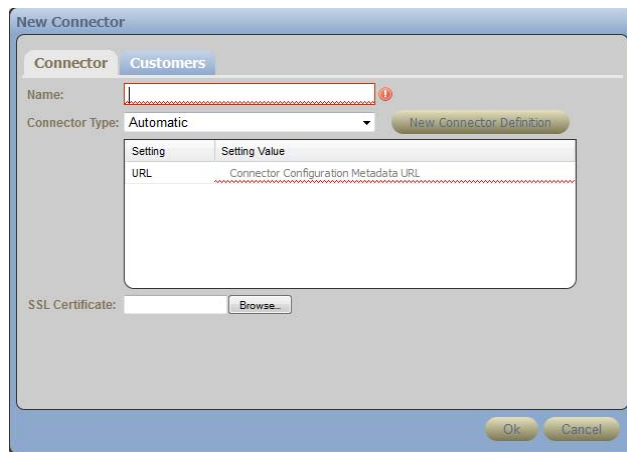
This section includes information about creating and configuring the NCSS connector for the Cloud Manager Application Console.

- ♦ “Adding a Connector for the Cloud Manager Application Console” on page 200
- ♦ “Adding Customers to the Cloud Manager Application Console Connector” on page 201
- ♦ “Configuring the Application Console Connector as a Customer” on page 202

Adding a Connector for the Cloud Manager Application Console

Use the steps in this section to add a surface connector in NCSS that interacts with Cloud Manager Application Console user interface. Before you follow these steps to configure a connector, make sure that Cloud Manager is running and configured for an [NCSS](#) or an [LDAP plus NCSS](#) authentication.

- 1 Log in to the NCSS Provider Console and open the Home page. For more information about logging in to the console, see “[Logging in to the Provider Console](https://www.netiq.com/documentation/cloudsecurityservice/provider/data/br69wb7.html)” (<https://www.netiq.com/documentation/cloudsecurityservice/provider/data/br69wb7.html>) in the *NetIQ Cloud Security Service 1.5 Provider Administration Guide*.
- 2 In the Connector panel, click *New Connector Definition*.



- 3 In the *Name* field, specify a name for this connector, identifying it as a connector to the Cloud Manager Application Console.
- 4 In the *Connector Type* field, select *Automatic*.
- 5 In the *URL* field, specify the full DNS name of your Cloud Manager Application Server, followed by the SSL port number, followed by `/config`. For example:
`https://netiq-office.dnsdhcp.city.netiq.com:8183/config`
- 6 In the *SSL Certificate* field, click *Browse*, then browse to and select the certificate [you created for the Cloud Manager Application Server](#). This certificate must be in DER format for NCSS, so you might need to create another copy converted from PEM format.
- 7 Click *OK* to create the new connector.

An error message similar to the following might be displayed:


```
Error Adding Connector, reason: ProxyException starting new Automatic connector
9c31b250-0ea6-4f88-98ac-9ab22f0df391; cause:
com.sun.jersey.api.client.ClientHandlerException:
javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.g:
PKIX path building failed: java.security.cert.CertPathBuilderException: unable
to find valid certification path to requested target.
```

The error displays for several minutes in the NCSS console before the connector eventually starts (that is, its icon changes from red to green). You can ignore the message content.

- 8 Continue with [“Adding Customers to the Cloud Manager Application Console Connector”](#) on page 201.

Adding Customers to the Cloud Manager Application Console Connector

After you have created the connector to the Cloud Manager Application Console, you need to establish which NCSS customers can access the Application Console.

- 1 On the *Home* page of the NCSS provider console, select the active (green) icon of the connector that you created in [Adding a Connector for the Cloud Manager Application Console](#) to display the Connectors page of the console.
- 2 On the *Connector Settings* panel of the *Connectors* page, click *Edit* to enable the edit mode for the *Customer List* subpanel.

Connector Settings [Save Changes] [Cancel]

Name:

SSL Certificate:

Customer List

Assigned Customers	Available Customers
democustomer	democustomer2

Provider-Specific Settings

Setting	Setting Value	Description
URL	https://ncss-demoncm.zenworks.lab.nov...	

Customer Setup

Customer-Specific Settings

Setting	Setting Value	Description
No Settings are required by this Connector.		

Mappings

Attribute	Description
Fdn	Fully-qualified LDAP distinguished name (AD: "distinguishedName")
Name	Name of user. Should be set from the user's "cn" LDAP attribute.
Roles	Cloud Manager group memberships (example: "groupmembership")
User Name	The Cloud Manager username. Should be set from the user's "mail" attribute.

Roles

Role	Description
No Roles are available for this Connector.	

- 3 Allow a customer access to the connector:
 - 3a In the *Available Customers* list, select the customer.
 - 3b Click the arrow icon to move the customer from the *Available Customers* list to the *Assigned Customers* list.
- 4 Click *Save Changes*.
- 5 Continue with [“Configuring the Application Console Connector as a Customer”](#) on page 202.

Configuring the Application Console Connector as a Customer

Cloud Manager requires that you map a customer's LDAP attributes to attributes defined within NCSS. This is necessary because eDirectory and Active Directory define their attributes in different ways. Without a common mapping, these divergent methods can cause problems for the Cloud Manager system. NCSS lets you map these attributes to a single, common set of definitions. This mapping is usually performed by the customer in the customer login view of NCSS, or you can do it as the NCSS provider administrator.

- 1 Log in to the NCSS console to open the Customers page. For more information, see *"Logging in to the Customer Console"* (<https://www.netiq.com/documentation/cloudsecurityservice/customer/data/br9x5xh.html>) in the *NetIQ Cloud Security Service 1.5 Customer Administration Guide* or *"Logging in to the Provider Console"* (<https://www.netiq.com/documentation/cloudsecurityservice/provider/data/br69wb7.html>) in the *NetIQ Cloud Security Service 1.5 Provider Administration Guide*.
- 2 (Conditional) If you log in to NCSS as a provider, you must select a customer in the *Customer List* panel of the Customers page. You must change the settings of the customer who will be logging in to the Cloud Manager Application Console.
- 3 In the *Connector List* panel of the Connectors page, make sure that the connector name of the Cloud Manager Application Console connector you created in [Adding a Connector for the Cloud Manager Application Console](#) is selected.
- 4 In the *Connector Settings* panel of the Connectors page, click *Edit*.
- 5 Map the attributes for the connector according to the LDAP source (eDirectory or Active Directory) you are using:
 - 5a Click the *Required Mappings* tab to open the required mappings table.
 - 5b Select the *Mapping Value* or *Type* column for each attribute. As you browse your LDAP source, consult the information in *Description* column to give you guidance, specify the values for each attribute, then click *Save*.
 - ♦ **FQDN:** Specify the fully-qualified LDAP distinguished name. In Active Directory, the value is the distinguishedName LDAP attribute. In eDirectory version 8.8.3 and later, the value is the entryDN LDAP attribute.
 - ♦ **Name:** Specify the name of this user. Obtain the value from the user's cn LDAP attribute.
 - ♦ **Roles:** Specify this user's group memberships in Cloud Manager. In eDirectory, this value is the groupMembership LDAP attribute. In Active Directory, this is the memberOf LDAP attribute.
 - ♦ **User Name:** Specify the Cloud Manager user name. Obtain the value from the user's mail attribute in his or her LDAP entry.
 - 5c In the *Connector Settings* panel, click *Save Changes*.
- 6 Select *Security Services* from the toolbar to open the *Identity Brokers* panel, then click *Send Configuration* to send the mappings to the NCSS Secure Bridge.
- 7 Continue with [Section 18.1.2, "Creating and Configuring an NCSS Connector for the Cloud Manager API,"](#) on page 203.

18.1.2 Creating and Configuring an NCSS Connector for the Cloud Manager API

To create and configure an NCSS connector for the Cloud Manager API, follow the steps in [Section 18.1.1, “Creating and Configuring an NCSS Connector for the Cloud Manager Application Console,” on page 200](#) that you used to [create the connector](#) (including its unique name), [associating users to the connector](#), and [configuring the connector as a customer](#).

The unique difference in creating a connector for the Cloud Manager API occurs in [Step 5](#) of the procedure, when you specify the full DNS name of your Cloud Manager Application Server in the *URL* field. For the API connector, use this syntax:

```
https://<ncm_server_address>:<port_if_needed>/cloudmanager-api/config
```

An actual URL might look like this:

```
https://netiq-office.dnshcp.city.netiq.com:8183/cloudmanager-api/config
```

When you have completed this process, continue with [Section 18.1.3, “Creating and Configuring an NCSS Connector for the Cloud Manager Tasks API,” on page 203](#).

18.1.3 Creating and Configuring an NCSS Connector for the Cloud Manager Tasks API

To create and configure an NCSS connector for the Cloud Manager Tasks API, follow the steps in [Section 18.1.1, “Creating and Configuring an NCSS Connector for the Cloud Manager Application Console,” on page 200](#) that you used to [create the connector](#) (including its unique name), [associating users to the connector](#), and [configuring the connector as a customer](#).

The unique difference in creating a connector for the Cloud Manager Tasks API occurs in [Step 5](#) of the procedure, when you specify the full DNS name of your Cloud Manager Application Server in the *URL* field. For the Tasks API connector, use this syntax:

```
https://<ncm_server_address>:<port_if_needed>/tasks-api/config
```

An actual URL might look like this:

```
https://netiq-office.dnshcp.city.netiq.com:8183/tasks-api/config
```

When you have completed this process, continue with [Section 18.1.4, “Creating and Configuring an NCSS Connector for the Cloud Manager App Services API,” on page 203](#).

18.1.4 Creating and Configuring an NCSS Connector for the Cloud Manager App Services API

To create and configure an NCSS connector for the Cloud Manager App Services API, follow the steps in [Section 18.1.1, “Creating and Configuring an NCSS Connector for the Cloud Manager Application Console,” on page 200](#) that you used to [create the connector](#) (including its unique name), [associating users to the connector](#), and [configuring the connector as a customer](#).

The unique difference in creating a connector for the Cloud Manager App Services API occurs in [Step 5](#) of the procedure, when you specify the full DNS name of your Cloud Manager Application Server in the *URL* field. For the App Services API connector, use this syntax:

```
/app-services-api/config
```

So an actual URL might look like this:

18.2 Configuring Novell Access Manager to Work with Cloud Manager

Novell Access Manager (NAM) provides secure, single sign-on access to trusted NetIQ Cloud Manager users from any location, in spite of the internal technical and organizational boundaries in your enterprise. Novell 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 Novell Access Manager and a Novell Access Manager Access Gateway. You should also have installed NetIQ Cloud Manager, and the Cloud Manager Application Server should be running.

You need to be familiar with Novell Access Manager capabilities so that you understand the context of the content in this section. For more information about Novell Access Manager, see the [Access Manager documentation Web site \(http://www.novell.com/documentation/novellaccessmanager31/index.html\)](http://www.novell.com/documentation/novellaccessmanager31/index.html).

- ♦ [Section 18.2.1, “Managing a Reverse Proxy for Authentication to Cloud Manager,” on page 204](#)

18.2.1 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”* (<http://www.novell.com/documentation/novellaccessmanager31/accessgatewayhelp/data/reverselist.html>) in the *Novell Access Manager 3.1 SP4 Configuration Guide*.

When the reverse proxy is set up as you want it, you need to perform the other configuration procedures necessary for Novell Access Manager authentication:

- ♦ [“Creating and Configuring the Proxy Service for the Cloud Manager Reverse Proxy” on page 205](#)
- ♦ [“Adding and Protecting All Cloud Manager Resources” on page 205](#)
- ♦ [“Creating an Identity Injection Policy for the New Cloud Manager Protected Resource” on page 207](#)
- ♦ [“Adding and Configuring an HTML Rewriter Profile for the Proxy Service” on page 209](#)

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* (<http://www.novell.com/documentation/novellaccessmanager31/accessgatewayhelp/data/b34l8ue.html>) in the *Novell Access Manager 3.1 SP4 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 205.


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”* (<http://www.novell.com/documentation/novellaccessmanager31/installation/data/b2euq6u.html>) in the *Novell Access Manager 3.1 SP4 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, make sure 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 207.

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”* (<http://www.novell.com/documentation/novellaccessmanager31/installation/data/b2euq6u.html>) in the *Novell Access Manager 3.1 SP4 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 209](#).

NOTE: Make sure that you always update your configuration when you make changes in Novell Access Manager.

For more information, see [“Configuring an Identity Injection Policy” \(http://www.novell.com/documentation/novellaccessmanager31/policyhelp/data/editpolicyii.html\)](http://www.novell.com/documentation/novellaccessmanager31/policyhelp/data/editpolicyii.html) in the *Novell Access Manager 3.1 SP4 Policy Guide*.


Adding and Configuring an HTML Rewriter Profile for the Proxy Service

The changes you make to the Novell 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 *“Understanding the Rewriting Process”* (<http://www.novell.com/documentation/novellaccessmanager31/accessgatewayhelp/data/b3nqotc.html#b3o4npk>) in the *Novell Access Manager 3.1 SP4 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”* (<http://www.novell.com/documentation/novellaccessmanager31/installation/data/b2euq6u.html>) in the *Novell Access Manager 3.1 SP4 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 box.



10 In the dialog box, specify the new content-type header, which is *application/xml*, select the *Rewrite Inbound Headers* check box, then click *OK* to make sure that the new Content-Type Header is enabled for the protected resource.

V Upgrading

PartIntro

- ♦ [Chapter 19, “Orchestration Components Upgrade Overview,” on page 213](#)
- ♦ [Chapter 20, “Upgrading Cloud Manager Orchestration Components,” on page 217](#)
- ♦ [Chapter 21, “Upgrading the Cloud Manager Application Server Components,” on page 233](#)
- ♦ [Appendix A, “Compatibility Checking Behavior for Orchestration Components,” on page 237](#)
- ♦ [Appendix B, “How to Recover from a Failed Orchestration Server Upgrade,” on page 241](#)

19 Orchestration Components Upgrade Overview

With this release, upgrade logic you implement for existing installations uses internal version numbers that might seem incongruent. NetIQ Cloud Manager Orchestration 3.1.4 components are upgraded to NetIQ Cloud Manager Orchestration Server 3.1.5 components, as follows:

- ♦ NetIQ Cloud Manager Orchestration Server 3.1.4 upgrades to NetIQ Cloud Manager Orchestration Server 3.1.5
- ♦ NetIQ Cloud Manager Orchestration Console 3.1.4 upgrades to NetIQ Cloud Manager Orchestration Console 3.1.5
- ♦ NetIQ Cloud Manager Orchestration Agent 3.1.4 upgrades to NetIQ Cloud Manager Orchestration Agent 3.1.5
- ♦ NetIQ Cloud Manager VM Builder 3.1.4 upgrades to NetIQ Cloud Manager VM Builder 3.1.5
- ♦ NetIQ Cloud Manager Monitoring Server 3.1.4 upgrades to NetIQ Cloud Manager Monitoring Server 3.1.5
- ♦ NetIQ Cloud Manager Monitoring Agent 3.1.4 upgrades to the NetIQ Cloud Manager Monitoring Agent 3.1.5

This section explains what you can expect from a NetIQ Cloud Manager Orchestration components 3.1.4 upgrade to NetIQ Cloud Manager Orchestration 3.1.5 components.

- ♦ [Section 19.1, “Basic Functions of the Orchestration Components Upgrade,” on page 213](#)
- ♦ [Section 19.2, “Cloud Manager Orchestration Components That Are Not Upgraded,” on page 214](#)

19.1 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, run the following command on a Linux machine where agent, client, or server components are installed.

```
rpm -qa | grep 'novell-zen'
```

To check version numbers 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 3.1.5 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 3.1.5 agents cannot communicate with the Cloud Manager Orchestration Server 3.1.4. You can upgrade the agents by selecting the Upgrade option on the Resource Registration dialog in the Cloud Manager 3.1.5 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, and the Cloud Manager VM Client 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 [Chapter A, “Compatibility Checking Behavior for Orchestration Components,” on page 237](#).
- ♦ After an upgrade to Cloud Manager Orchestration Server 3.1.5, 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 B.1, “Upgrade Failure Scenarios,” on page 241](#).
- ♦ After the Cloud Manager Orchestration Server 3.1.4 is upgraded to Cloud Manager Orchestration Server 3.1.5, rolling back to Cloud Manager Orchestration Server 3.1.4 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 20, “Upgrading Cloud Manager Orchestration Components,” on page 217](#).

19.2 Cloud Manager Orchestration Components That Are Not Upgraded

When you upgrade from Cloud Manager Orchestration 3.1.4 to Cloud Manager Orchestration 3.1.5, 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.

For example, suppose you have deployed the xen provisioning adapter job and you made custom changes to the xen policy file. When the Orchestration Server prepares for an upgrade, it repackages the xen provisioning adapter by creating a `.sar` archive and then stores it in `/Orchestration_instance_directory/snapshot/deployment/core/xen.sar`. This `xen.sar` archive contains the current state of the xen provisioning adapter, including your custom changes.

Later, when the Orchestration Server is upgraded, the new xen provisioning adapter for the new server is deployed, but the changes you made previously are not applied. To apply these changes to the new server, you have two choices:

- ♦ Use the Cloud Manager Orchestration Console to manually apply the changes to the new server's core component. (You can review what these changes were by looking at the snapshot files in the `xen.sar` archive.)
- ♦ (Conditional) If you are migrating between servers of the same version whose core components have not changed, you can use the `zosadmin redeploy` command to manually redeploy the snapshotted core component.

NOTE: After the upgrade to Cloud Manager Orchestration 3.1.5 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.

20 Upgrading Cloud Manager Orchestration Components

This section provides information about upgrading from earlier NetIQ Cloud Manager Orchestration components to current NetIQ Cloud Manager Orchestration components. It is important that you upgrade the Orchestration components you have installed in the sequence that follows:

- ♦ [Section 20.1, “Upgrading Orchestration Components,” on page 217](#)
- ♦ [Section 20.2, “Alternate Methods for Upgrading Older Agents and Clients,” on page 229](#)
- ♦ [Section 20.3, “Running the Upgrade Configuration on an Enterprise Scale,” on page 231](#)
- ♦ [Section 20.4, “Upgrading a Cloud Manager Orchestration High Availability Configuration,” on page 232](#)

20.1 Upgrading Orchestration Components

The following information lists the upgrade steps in the order that they should be performed.

1. [Section 20.1.1, “Backing Up the Orchestration Components Prior to Upgrading,” on page 218](#)
2. [Section 20.1.3, “Checking the Current Version of Cloud Manager Orchestration Components,” on page 219](#)
3. [Section 20.1.4, “Snapshotting the Existing Orchestration Server Installation,” on page 219](#)
4. [Section 20.1.5, “Upgrading the Orchestration Packages,” on page 220](#)
5. [Section 20.1.6, “Checking the Upgraded Version of the Orchestration Components,” on page 223](#)
6. [Section 20.1.7, “Configuring the Upgraded Packages,” on page 224](#)
7. [Section 20.1.8, “Manually Configuring the Remote Audit Database after Orchestration Components Are Upgraded,” on page 227](#)
8. [Section 20.1.9, “Upgrading the XenServer Provisioning Adapter,” on page 228](#)
9. [Section 20.1.10, “Running Discovery on VM Hosts and Images,” on page 228](#)

NOTE: To perform a mass upgrade of NetIQ 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 example, you can use Novell ZENworks Linux Management to distribute new agents and clients to Linux servers.

If you choose to use ZENworks Linux Management, you should enable the rollback command. This will let you easily roll back to the prior version of Cloud Manager Orchestration if the upgrade to newer Orchestration components is unsuccessful.

For more information, see [Section 20.3, “Running the Upgrade Configuration on an Enterprise Scale,” on page 231](#).

20.1.1 Backing Up the 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 version:

1. Make a copy of the directories under `/var/opt/novell/zenworks/zos/`
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.

20.1.2 Backing Up the Application Components Prior to Upgrading

Anytime before you upgrade, you can back up the application components in the Cloud Manager 2.1.4 system, including any custom files (with the exception of the Postgres database). Cloud Manager 2.1.4 includes a shell script tool that lets you back up the current Cloud Manager 2.1.3 system prior to the upgrade so that if the upgrade to 2.1.4 fails, you can revert to the 2.1.3 configuration settings and files.

NOTE: A reversion to Cloud Manager 2.1.3 settings and files pre-supposes that you would revert to an earlier version of the Postgres database.

Use these steps to perform the pre-upgrade backup:

- 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 2.1.3 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 2.1.3 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 2.1.3 system like this:

```
./backup -r -f /var/opt/netiq/cloudmanager/my_backup_name.tar.gz
```

20.1.3 Checking the Current Version of Cloud Manager Orchestration Components

Before you upgrade the Cloud Manager Orchestration packages from version 3.1.5 to the Cloud Manager Orchestration 3.2.0 packages, you should check which 3.1.5 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 20.1.6, “Checking the Upgraded Version of the Orchestration Components,” on page 223](#)).

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.

20.1.4 Snapshotting the Existing Orchestration Server Installation

Before you begin the upgrade process of the Orchestration Server, make sure that all running jobs are complete. 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`.

Use the following steps 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 box.



Select *Perform Snapshot of Server State*, then click *Shutdown*.

20.1.5 Upgrading the Orchestration Packages

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 220](#).
- If you want to use the command line to upgrade, see [“Upgrading Orchestration Packages Using the zypper Command” on page 222](#).
- If you use ZENworks Linux Management tools to upgrade the packages, we recommend that you use the same tools to clean up the environment, see the [ZENworks Linux Management documentation Web site \(http://www.novell.com/documentation/zlm72/\)](http://www.novell.com/documentation/zlm72/) for more information.

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 2.2.0 ISO (64-bit), then prepare it for installation:

TIP: The NetIQ Cloud Manager 2.2.0 product ISO includes the packages for Cloud Manager Orchestration components, which carry the 3.2.0 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-2.2.0-<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 RPMs are upgraded, scripts are run that do the following:
- ♦ Back up the existing server instance directory
 - ♦ Upgrade the RPMs 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 Packages” on page 224](#).
 - ♦ Run the GUI Configuration Wizard. If you use this method, skip to the steps in [“Configuring the Upgraded Packages” on page 224](#).

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 11x machines. If you want to use the GUI Configuration Wizard to upgrade, see [“Configuring an Upgraded Orchestration Agent Installed on the Server” on page 227](#).

For more `zypper` commands, see [“Other Useful zypper Commands for Upgrade” on page 223](#).

- 1** Download the appropriate Cloud Manager 2.1.4 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 2.1.3 to Cloud Manager 2.1.4:
 - 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-2.1.4-SLE11.x86_64.iso, you could use this command:

```
zypper sa -t yast2 "iso:/?iso=/root/Desktop/NetIQ_Cloud_Manager-2.1.4-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>
```

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 20-1 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.

20.1.6 Checking the Upgraded Version of the Orchestration Components

After you upgrade the Cloud Manager Orchestration 3.1.5 packages to Cloud Manager Orchestration 3.2.0 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 20.1.3, “Checking the Current Version of Cloud Manager Orchestration Components,”](#) on page 219). If some of the components have not been upgraded from the earlier version, the incompatibility between the components could cause unexpected behavior.

20.1.7 Configuring the Upgraded Packages

This section discusses the basic upgrade configuration of all NetIQ 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:

- ♦ [“Some Considerations When Configuring with the GUI Wizard” on page 224](#)
- ♦ [“Configuring the Upgraded Orchestration Server” on page 224](#)
- ♦ [“Configuring the Upgraded Monitoring Server” on page 227](#)
- ♦ [“Configuring an Upgraded Orchestration Agent Installed on the Server” on page 227](#)

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 Upgrade of the Orchestration Server” on page 226](#).

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 Cloud Manager component.

- 1 Make sure 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 <i>u</i> (for upgrade) in the configuration script, or select the <i>Upgrade</i> check box in the wizard.

Server Upgrade Configuration Requirement	Explanation and Action
Configuration Type	<p>Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation.</p> <p>This section discusses standard installation, so specify <code>s</code> (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 153.</p>
Administrator User	IMPORTANT: Do not change the administrator name from the name you used in the original installation.
Administrator Password	<p>Enter the password of the Cloud Manager Orchestration Server administrator you used in the previous installation (for Cloud Manager 2.1.5/Orchestration Server 3.1.5).</p> <p>IMPORTANT: Do not change the administrator password from the password you used in the original installation.</p>
Confirm Password / Retype Password	Re-enter the administrator password.
Automatic Agent Upgrade	<p>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.</p> <p>You can enable automatic upgrade of old agents to avoid this approval step if you enter <code>yes</code> or check the <i>Automatic Agent Upgrade</i> check box (wizard).</p> <p>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 “Server Admin Object” section of the <i>NetIQ Cloud Manager Component Reference</i>.</p>
Upgrade Auditing Database	<p>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.</p> <p>If you use a different RDBMS, you need configure it separately.</p> <p>If you select the <i>Upgrade Audit Database</i> check box (wizard), or if you enter <code>yes</code> for this prompt, you must specify</p> <ul style="list-style-type: none"> ♦ 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 you 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. <p>If you want to manually configure the database after upgrade, see Section 20.1.8, “Manually Configuring the Remote Audit Database after Orchestration Components Are Upgraded,” on page 227.</p>

Server Upgrade Configuration Requirement	Explanation and Action
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	<p>As you upgrade, you need a license key (90-day evaluation license or a full license) to use this product.</p> <p>Depending on your arrangements with NetIQ, you either received a new key or you were given permission to reuse your old key.</p> <p>Specify the path to the network location of either your existing key or the new key that was issued to you.</p>

- 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 Upgrade 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 224](#).

- 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 3.1.5 to the Orchestration Server 3.2.0.

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.

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.

20.1.8 Manually Configuring the Remote Audit Database after Orchestration Components Are Upgraded

When 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
```

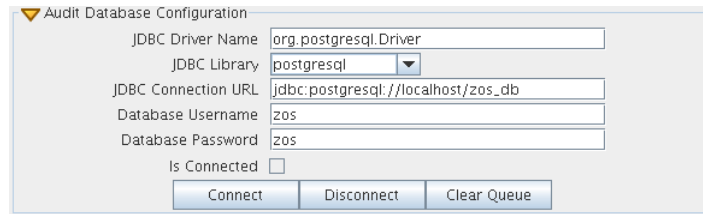
```
psql -h <psql-server-addr> -d postgres -U postgres -f audit_db_upgrade.sql
```

- 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 username and password match the values used when creating the schema owner database user in [Section 7.1, “Configuring the Orchestration Server,”](#) on page 64. In this example, the username is `zos` and the password is `zos`.



- 5 Confirm that the database username and password match the values you replaced in the variables of the `.sql` script. In this example, the username 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.

20.1.9 Upgrading the XenServer Provisioning Adapter

When upgrading the orchestration components from version 3.1.3 to 3.1.4, the XenServer provisioning adapter needs to be handled separately. To upgrade this provisioning adapter, perform the following steps:

- 1 Shutdown and remove the existing agent on the XenServer hosts.

- 1a Connect to the XenServer host

- 1b Shut down the Orchestrator Agent:

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

- 1c Delete the RPMs installed to support the old provisioning adapter and the Orchestrator Agent:

```
rpm -e cabextract libmspack chntpw fuse-ntfs-3g fuse novell-zenworks-zos-agent novell-zenworks-zos-java
```

- 1d Remove any remaining files from the previous installation:

```
rm -rf /opt/novell
rm -rf /var/opt/novell
```

- 2 Perform a fresh installation of the XenServer provisioning adapter as outlined in [Section 11.2, “Configuring the Citrix XenServer Provisioning Adapter,”](#) on page 95.
- 3 Continue with [Section 20.1.10, “Running Discovery on VM Hosts and Images,”](#) on page 228.

20.1.10 Running Discovery on VM Hosts and Images

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*.

20.2 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 3.2.0 components were installed. This section includes information that helps you to walk through the upgrade of those agents and clients.

- ♦ [Section 20.2.1, “Automatically Upgrading the Orchestration Agent from the Cloud Manager Orchestration Console,” on page 229](#)
- ♦ [Section 20.2.2, “Using the ISO to Upgrade the Orchestration Agent on Red Hat Enterprise Linux 5 Machines,” on page 230](#)
- ♦ [Section 20.2.3, “Using the ISO to Upgrade the Old Orchestration Agent or the Orchestration Clients on Windows Machines,” on page 231](#)
- ♦ [Section 20.2.4, “Using the Administrator Information Page to Upgrade the Agents and Clients,” on page 231](#)

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 2.2.0. For example, you can use ZENworks Linux Management to distribute new agents and clients to Linux servers.

For more information, see [Section 20.3, “Running the Upgrade Configuration on an Enterprise Scale,” on page 231](#).

20.2.1 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 2.2.0) 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 box displays all of the older resource objects attempting to connect. The *Upgrade* option is also available in the dialog box. 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`
- ♦ `cabextract` (on Xen hosts where you plan to use Windows sysprep)

To separately upgrade these components, you need to use the YaST upgrade from the product ISO (see [“Upgrading the Orchestration Packages” on page 220](#)) 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 222](#).

20.2.2 Using the ISO to Upgrade the Orchestration Agent on Red Hat Enterprise Linux 5 Machines

Use the following procedure if you want to use the Add-on method to upgrade the Cloud Manager 3.1.5 Orchestration Console to a Cloud Manager 3.2.0 Orchestration Console running on a Red Hat Enterprise Linux (RHEL) 5 machine.

- 1 Shut down the old Orchestration Console on the machine where you intend to install the new Cloud Manager 3.2.0 Orchestration Console.

- 2 Download the appropriate Cloud Manager 2.2.0 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-2.2.0-SLE11.x86_64.iso /mnt
```

This mounts the ISO in the /mnt folder.

- 4 Navigate to the directory path where the RHEL 5 packages reside. For example:

```
cd /mnt/RHEL5
```

There are four packages in the /RHEL5 directory:

```
netiq-ncm-orch-config-3.1.5-<build>.noarch.rpm
```

```
netiq-ncm-orch-config-gui-3.1.5-<build>.noarch.rpm
```

```
netiq-ncm-cmos-agent-3.1.5-<build>.x86_64.rpm
```

```
netiq-ncm-cmos-java-1.6.0_sun_update14-1.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 `libcap1so` 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.

NOTE: Upgraded agent and client software does not require you to execute the configuration script on RHEL 5 machines.

20.2.3 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 Windows XP, Windows Vista and Windows 7 desktops. To upgrade, install the new Cloud Manager 2.2.0 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 (Windows XP, Windows Vista, or Windows 7) reside.
- 4 Double-click the appropriate file (.exe) to launch an installation and configuration wizard for the Orchestration Console.

20.2.4 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 45](#)). 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 upgrading the Cloud Manager Orchestration Console on a Windows machine, you must uninstall the prior version first, then install the new version.

20.3 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.

20.4 Upgrading a Cloud Manager 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 20.1.4, “Snapshotting the Existing Orchestration Server Installation,”](#) on page 219 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 RPMs on the first node (the one you started in [Step 3](#)) of the cluster. For more information, see [Section 20.1.5, “Upgrading the Orchestration Packages,”](#) on page 220.
- 7 Run the configuration script. For more information, see [“Configuring the Upgraded Packages”](#) on page 224.
- 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 RPMs on the second node of the cluster.
- 14 Using the high availability manager (see [Step 1](#)), start the Orchestration Server instance.

21 Upgrading the Cloud Manager Application Server Components

This section includes information you need for upgrading the NetIQ Cloud Manager Application Server 2.1.5 to Cloud Manager Application Server 2.2.0. It also includes information about restoring the Cloud Manager Application Server system in the event of a disaster.

- ♦ [Section 21.1, “Backing Up the PostgreSQL Database,” on page 233](#)
- ♦ [Section 21.2, “Performing a Complete Cloud Manager System Backup,” on page 234](#)
- ♦ [Section 21.3, “Running the Cloud Manager Configuration Script,” on page 234](#)
- ♦ [Section 21.4, “Upgrading from a Pre-2.1.5 Version of Cloud Manager,” on page 235](#)
- ♦ [Section 21.5, “Restoring Cloud Manager In the Event of a System Failure,” on page 236](#)

For information about upgrading the Cloud Manager Orchestration Server, see [Chapter 20, “Upgrading Cloud Manager Orchestration Components,” on page 217](#).

21.1 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](#). The following specific instructions are helpful in this process:

- ♦ If you chose to automatically [configure the Postgres database](#) during the initial Cloud Manager installation, the Postgres password is located in the `/etc/opt/netiq/cloudmanager/etc/pgusr.in` file.
- ♦ If the Postgres 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
```

21.2 Performing a Complete Cloud Manager System Backup

If you want to perform a full Cloud Manager system backup (excluding the Orchestration components), you can do so at any time prior to the upgrade process. This backup saves all Cloud Manager files and any custom files you might have created for it. This is an optional process; normally, the backup of the Cloud Manager configuration automatically triggered during an upgrade is sufficient for Cloud Manager 2.1.5 restoration.

For more information about performing a full backup, see [Appendix 20.1.2, “Backing Up the Application Components Prior to Upgrading,” on page 218](#).

21.3 Running the Cloud Manager Configuration Script

When you have successfully installed the new Cloud Manager 2.2.0 package, you can 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.

Use the following procedure to run the configuration program:

- 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 Postgres user name that you specified in the previous Cloud Manager installation.
- 4 Specify the password for the Postgres administrator, then press Enter.
This must be the same Postgres password that you specified in the previous Cloud Manager installation.
- 5 (Conditional, if you have backed up the Postgres database, as explained in [“Backing Up the PostgreSQL Database” on page 233](#)). Specify yes to indicate that you have backed up the Postgres database, then press Enter.
- 6 After the configuration summary is displayed, specify yes to start the configuration process.

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`.

21.3.1 Using the Configurator Tool to Update Resource Pool Data on the Cloud Manager Application Server

The Cloud Manager configurator utility performs a VMware resource pool ID data upgrade after all Orchestration grids (also known as Cloud Manager Zones) are upgraded. 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.

21.4 Upgrading from a Pre-2.1.5 Version of Cloud Manager

If you are upgrading from a version of Cloud Manager earlier than 2.1.5 and if your current Cloud Manager environment has resource groups that include repositories, it is necessary to execute a command at the Karaf console of the Cloud Manager Application Server after you upgrade to version 2.2.0 and before you start the upgraded workloads.

```
cm:repo-groups-upgrade
```

The command executes a number of operations:

- Ensures that the `NCM_RGRepository` policy is deployed on the Orchestration Server.
- Ensures that the repository group is created on the Cloud Manager 2.2.0 Orchestration Server.
- Sets the `resource.vm.ncm.repository_group` fact on the VMs that are associated with a resource group.

You should also run the command when you create a new workload in an upgraded business service or when you import a VM and associate it with an upgraded resource group.

21.5 Restoring Cloud Manager In the Event of a System Failure

Should a failure happen to occur in your Cloud Manager 2.2.0 system, you can revert to Cloud Manager 2.1.5 by using the following steps:

- 1 Ensure that all Cloud Manager services are shut down. Use the following command at the Application Server command line:
- 2 Uninstall the Cloud Manager 2.2.0 RPMs. Use the following command at the Application Server command line:

```
/etc/init.d/netiq/cloudmanager clean
```

```
rpm --erase netiq-cloudmanager-2.2.0 netiq-cloudmanager-lib-2.2.0
```

- 3 Install the Cloud Manager 2.1.5 RPMs. Use the following command at the Application Server command line:

```
rpm -Uvh netiq-cloudmanager-2.1.5-<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-<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_2.2.0>.tar.gz
```

- 5 Restore the [backup of the Postgres](#) database you used before the upgrade.

IMPORTANT: The database you have been using with Cloud Manager 2.2.0 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
```

A 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.

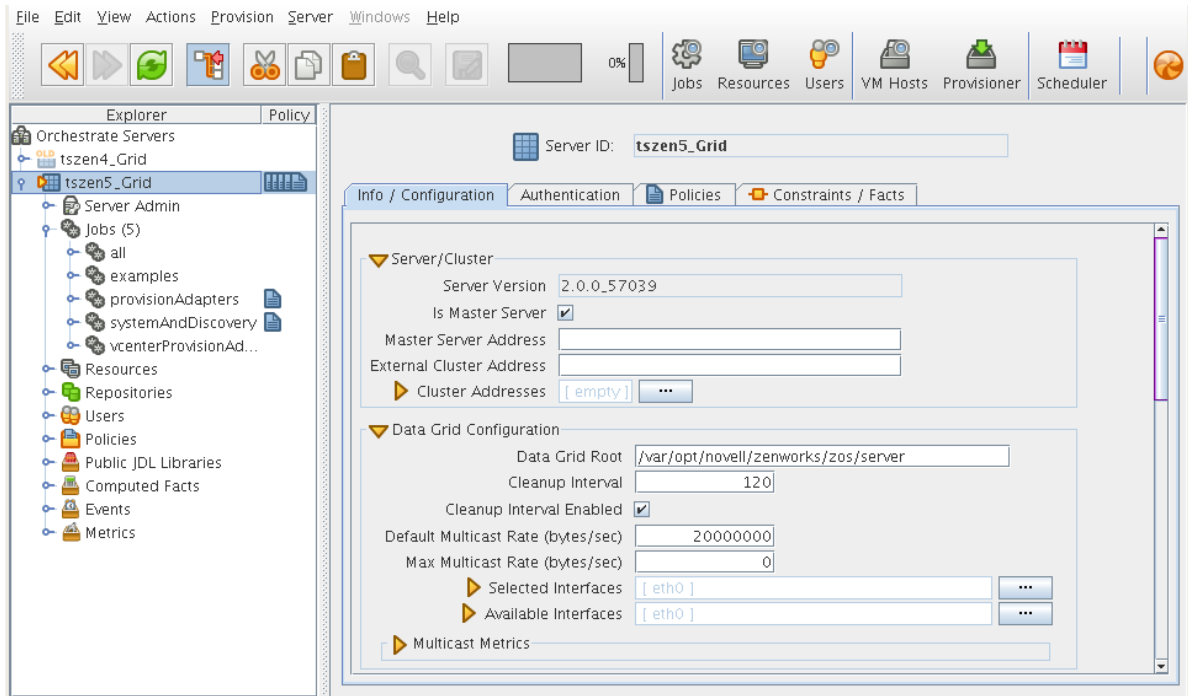
The Orchestration Console detects incompatibility only in the Orchestration Agent. The information in the following sections details that behavior.

- ♦ [Section A.1, “If the Orchestration Server Is Not Compatible with the Orchestration Console,” on page 237](#)
- ♦ [Section A.2, “When an Agent Version Does Not Match the Server Version,” on page 238](#)

A.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 A-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.

A.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 A-2 Old Orchestration Agent Resource Displayed in Tree View





Figure A-3 Old Orchestration Agent Resource Displayed in Tree View and Monitor View



B How to Recover from a Failed Orchestration Server Upgrade

The information in this section can help you to recover from a failed Orchestration Server upgrade.

- ♦ [Section B.1, “Upgrade Failure Scenarios,” on page 241](#)
- ♦ [Section B.2, “Restoring the Orchestration Server If the Upgrade Fails,” on page 242](#)

B.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 B.1.1, “Failure Scenario 1: Error Resolution,” on page 241](#)
- ♦ [Section B.1.2, “Failure Scenario 2: Cannot Resolve Error,” on page 241](#)

B.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

B.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 3.1.5 data from the snapshot.
- 3 Restore the previous version RPMs of the Orchestration Server 3.1.5 software.

B.2 Restoring the Orchestration Server If the Upgrade Fails

If you use ZENworks Linux Management in your network, you can use it to restore an older version of the Cloud Manager Orchestration Server if an upgrade has failed. This section contains information that can help you roll back a failed upgrade of Cloud Manager Orchestration Server 3.2.0 back to Cloud Manager Orchestration Server 3.1.5.

- ♦ [Section B.2.1, “Requirements,” on page 242](#)
- ♦ [Section B.2.2, “Rollback Procedure Using the rug Command,” on page 242](#)

B.2.1 Requirements

This scenario requires that you have already installed the Cloud Manager 2.1.5 ISO. That is, the 3.1.5 version of the Orchestration Server should be running with exactly the same packages you originally installed and configured.

The scenario also requires that you have a Cloud Manager 2.2.0 ISO on hand. It is important that you enable rollback through ZENworks Linux Management before you actually execute the rollback. ZENworks Linux Management records the changes you make to the RPM database when you enable rollback.

Rollback works only if you previously installed 3.1.5 packages using ZENworks Linux Management. ZENworks Linux Management records data about each package that it installs, deletes, or upgrades.

For more information about using ZENworks Linux Management for rollback, see [Reverting to a Previously Installed Software Configuration State \(http://www.novell.com/documentation/zlm73/lm7admin/data/b94fftd.html\)](http://www.novell.com/documentation/zlm73/lm7admin/data/b94fftd.html) in the *ZENworks 7.3 Linux Management Administration Guide*.

B.2.2 Rollback Procedure Using the rug Command

Use the following steps to roll back a Cloud Manager Orchestration Server 3.2.0 upgrade to Cloud Manager Orchestration Server 3.1.5 on SLES 11x machines where the ZENworks Linux Management Daemon is also installed.

- 1 Use the following command to make sure that you have the ZENworks Management Daemon installed, with rollback tools enabled.

```
rug get rollback
```

- 2 Check repositories to ensure that they are disabled. You want only 3.2.0 upgrades.

- 2a Run the following command to list the repositories:

```
rug sl
```

- 2b Run the following command to list the catalogs of subscribed repositories:

```
rug ca
```

- 2c Run the following command to unsubscribe from each subscribed repository:

```
rug unsub "<name_of_repository>"
```

- 3 Add a Cloud Manager Orchestration Server 3.1.5 ISO as a repository.

- 3a Run the following command, followed by the local path of the ISO, the ftp or http addresses, or the path to the CD or DVD media where the installation source of Cloud Manager Orchestration Server 3.1.5 currently resides.

```
rug sa -t zypp <installation_source_of_Orchestration_Server_3.2.0> cmos315
```

This command adds the Orchestration Server 3.1.5 repository to the ZENworks Management Daemon. The daemon uses the RPMs in the repository to roll back the server to its former state. For this reason, the repository (the `cmos315` shown in the example) must have the same RPM package versions as Orchestration Server 3.1.5.

For more information about adding repositories, see the [ZENworks 7.3 Linux Management Administration Guide](http://www.novell.com/documentation/zlm73/lm7admin/data/front.html) (<http://www.novell.com/documentation/zlm73/lm7admin/data/front.html>).

- 3b** Run the following command to list and confirm existing repositories:

```
rug sl
```

- 3c** Run the following command to list and confirm the catalogs of subscribed repositories:

```
rug ca
```

- 4** Subscribe to the `cmos315` repository.

- 4a** Run the following command to subscribe to the `cmos315` repository:

```
rug sub cmos315
```

- 4b** Run the following command to list the catalogs and confirm the catalogs of subscribed repositories:

```
rug ca
```

The new repository shows `Yes` in the `Sub'd` (subscribed) column.

- 4c** Run the following command to list and confirm updates:

```
rug lu
```

The message, `No updates are available`, is displayed, which indicates that no new updates to the repository are available—the RPMs match those in the `ncm20` catalog.

- 5** Add a Cloud Manager 2.2.0 ISO as a repository.

- 5a** Run the following command, followed by the local path of the ISO, the ftp or http addresses, or the path to the CD or DVD media where the installation source of Cloud Manager 2.2.0 currently resides.

```
rug sa -t zypp <installation_source_of_Cloud  
Manager_2.2.0_Orchestration_Server> cmos320
```

This command adds the Orchestration Server 3.2.0 repository to the ZENworks Management Daemon. The daemon uses the RPMs in the repository to roll back the server to its former state. For this reason, the repository (the `cmos320` shown in the example) must have the same RPM package versions as Orchestration Server 3.2.0.

For more information about adding repositories, see the [ZENworks 7.3 Linux Management Administration Guide](http://www.novell.com/documentation/zlm73/lm7admin/data/front.html) (<http://www.novell.com/documentation/zlm73/lm7admin/data/front.html>).

- 5b** Run the following command to list and confirm existing repositories:

```
rug sl
```

- 5c** Run the following command to list and confirm the catalogs of subscribed repositories:

```
rug ca
```

- 6** Subscribe to the `cmos320` repository.

- 6a** Run the following command to subscribe to the `cmos320` repository:

```
rug sub cmos320
```

- 6b** Run the following command to list the catalogs and confirm the catalogs of subscribed repositories:

```
rug ca
```

The new repository shows **Yes** in the **Sub'd** (subscribed) column for both the **cmos315** and **cmos320** repositories.

- 6c** Run the following command to list and confirm the updated Orchestration Server 3.1.x packages:

```
rug lu
```

- 7** Run the following command to verify that the 3.1.5 Server and the 3.1.5 Agent are in a running state:

```
ps ax | grep java
```

- 8** Run the following command to perform the package upgrade (while the server is in a running state).

```
rug up
```

The upgrade scripts of the Orchestration Server 3.2.0 RPM packages stop the Orchestration Server 3.1.5 and the 3.1.5 Agent before the upgrade, then take a snapshot of the 3.1.5 Server that is required for the upgrade.

- 9** When the package upgrade is complete, run the following command to launch the configuration script to upgrade the 3.1.5 Server.

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

NOTE: For details on running the configuration script, see [Chapter 7, “Configuring Cloud Manager Orchestration Components,”](#) on page 63.

If the upgrade configuration fails, error information is displayed in the terminal.

Because of the configuration upgrade failure, you need to use the ZENworks Management Daemon to roll back to the former (that is, the Orchestration Server 3.1.5) running state without losing data.

- 10** Run the following command to confirm that Cloud Manager Orchestration Server 3.2.0 packages are installed:

```
rpm -qa | grep zos
```

Because the Orchestration Server 3.2.0 packages are installed but not configured, you cannot use them to start the Orchestration Server.

- 11** Run the follow command to confirm that an instance of the Orchestration Server was created:

```
ls /var/opt/novell/zenworks/zos/
```

The `/agent`, `/server` and `/server.save` folders should be listed.

- 12** Run the following command to launch the ZENworks Linux Management (that is, the ZENworks Management Daemon) for rolling back to Orchestration Server 3.1.5:

```
rug ro 1 hour ago
```

NOTE: The rollback parameter, `1 hour ago`, is conditional: it specifies the state of the packages on the SLES server at a given time in the past. You need to specify the time when you are sure that 3.1.5 packages were installed and running so that you can roll back the current Orchestration Server 3.2.0 packages to Orchestration Server 3.1.5 packages.

- 13** Run the following commands to confirm that the system has been rolled back to version 3.1.5 and that a server instance exists:

```
rpm -qa | grep zos
```

```
ls /var/opt/novell/zenworks/zos/
```

14 Run the following command to start the Orchestration Agent:

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

VI Uninstalling

Part Intro

- ♦ [Chapter 22, “Uninstalling Orchestration Component Patterns from a SLES Server,”](#) on page 249

22 Uninstalling Orchestration Component Patterns from a SLES Server

This section includes information about uninstalling NetIQ 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 YaST and YaST2 is not 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 is a variety of tools you can use to perform the uninstall:

- ♦ Use the `rug` command line tool on SUSE Linux Enterprise Server (SLES) 10 machines to remove the server packages.

```
rug remove -t pattern <pattern_name>
```

- ♦ Use the `zypper` command line tool to remove the server packages on SLES 11 machines.

```
zypper remove <RPM_Package>
```
- ♦ Use the YaST setup and configuration tool on either SLES 10 or SLES 11 machines to select and remove the software packages.
- ♦ Use the `rpm -qa` command to look for installed RPMs 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 RPMs:
 - ♦ `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

23 What's Next?

When you complete installation and configuration of your NetIQ Cloud Manager system, as explained in this book, you are ready to start populating your system with the components that enable users to provision their own business services. For information, see the [NetIQ Cloud Manager Procedures Guide](#).

