Orchestration Server High Availability Configuration Guide

Cloud Manager 2.1.4

November 14, 2012

Legal Notice

Contents

About This Guide

This *High Availability Configuration Guide* provides the information for installing and configuring NetIQ Cloud Manager Orchestration Server in a high availability environment. The guide provides information about the components and configuration steps necessary for preparing this environment, including instructions for configuring the Orchestration Server in a cluster. The guide also provides some information regarding the behaviors you can expect from the Orchestration Server in various failover scenarios. The guide is organized as follows:

Intended Audience

This information is intended for anyone who is assigned the Cloud Administrator role for a NetIQ Cloud Manager system. Consumers of this information should be experienced Linux and Windows system administrators who are familiar with virtual machine technology and datacenter operations.

Additional Documentation

For other NetIQ Cloud Manager 2.1.4 documentation, see the NetIQ Cloud Manager 2.*x* documentation site (https://www.netiq.com/documentation/cloudmanager2/).

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. We want to hear your comments and suggestions about this manual and the other documentation included with this product.

- Please use the *User Comments* feature at the bottom of each page of the online documentation to provide specific feedback about the content on that page. A documentation representative will contact you via e-mail with a resolution to the documentation problem within five business days.
- If you have more general suggestions for improvements, please 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.

1 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.

1.1 Overview

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

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

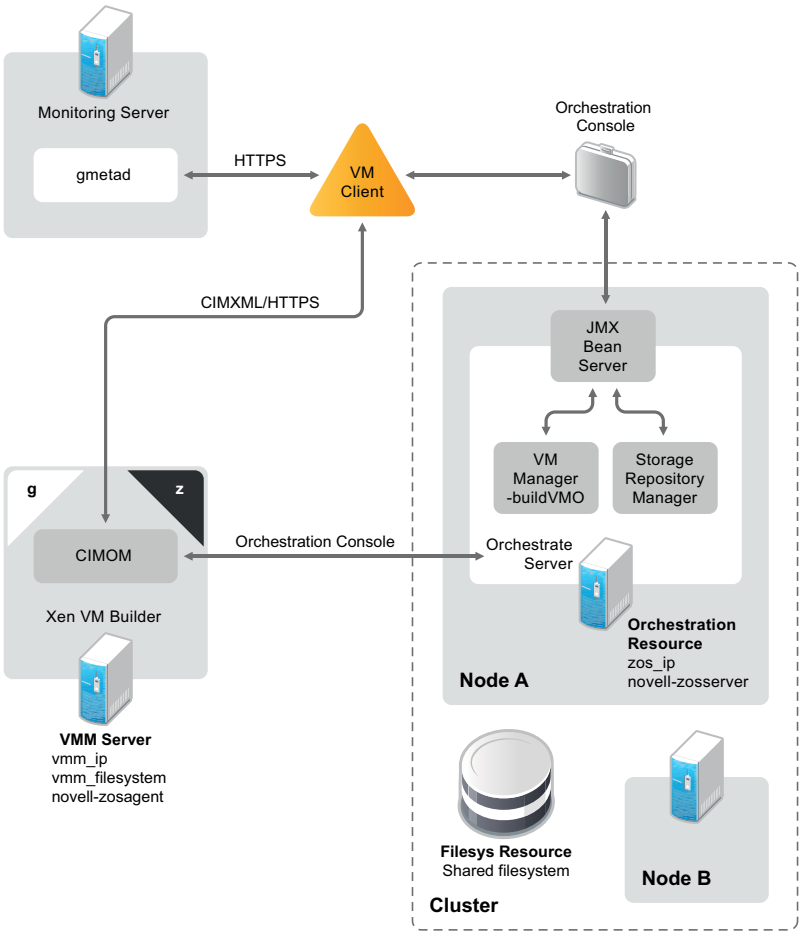

1.2 Installing the Orchestration Server to a SLES 11 Pacemaker Cluster Environment

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 1.2.1, "Meeting the Prerequisites," on page 11
2. Section 1.2.2, "Installing the SLES 11 SP2 High Availability Pattern," on page 12
3. Section 1.2.3, "Configuring SLES 11 Nodes with Time Synchronization and Installing Pacemaker to Each Node," on page 13
4. Section 1.2.4, "Setting Up OCFS2 on SLES 11 SP2," on page 14
5. Section 1.2.5, "Installing the Orchestration Server on the First Clustered SLES 11 Node," on page 15

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 the *NetIQ Cloud Manager SUSE Xen VM High Availability Configuration Guide*.

NOTE: Upgrading from earlier versions of Cloud Manager Orchestration to a high availability environment is supported. For more information, see "Upgrading Cloud Manager Orchestration 3.1.3 Components to Cloud Manager Orchestration 3.1.4 Components" in the *NetIQ Cloud Manager 2.1.4 Orchestration Components Upgrade Guide*.

1.2.1 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 11
- "Software Requirements for Creating a High Availability Environment" on page 11

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=9xvsJDAsS04~).

  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 1.2.4, "Setting Up OCFS2 on SLES 11 SP2," on page 14 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.

1.2.2 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*.

1.2.3 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 13
- "Configuring Pacemaker" on page 13

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) 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) 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 1-2 *The Main Settings Page in the Pacemaker Graphical Interface*

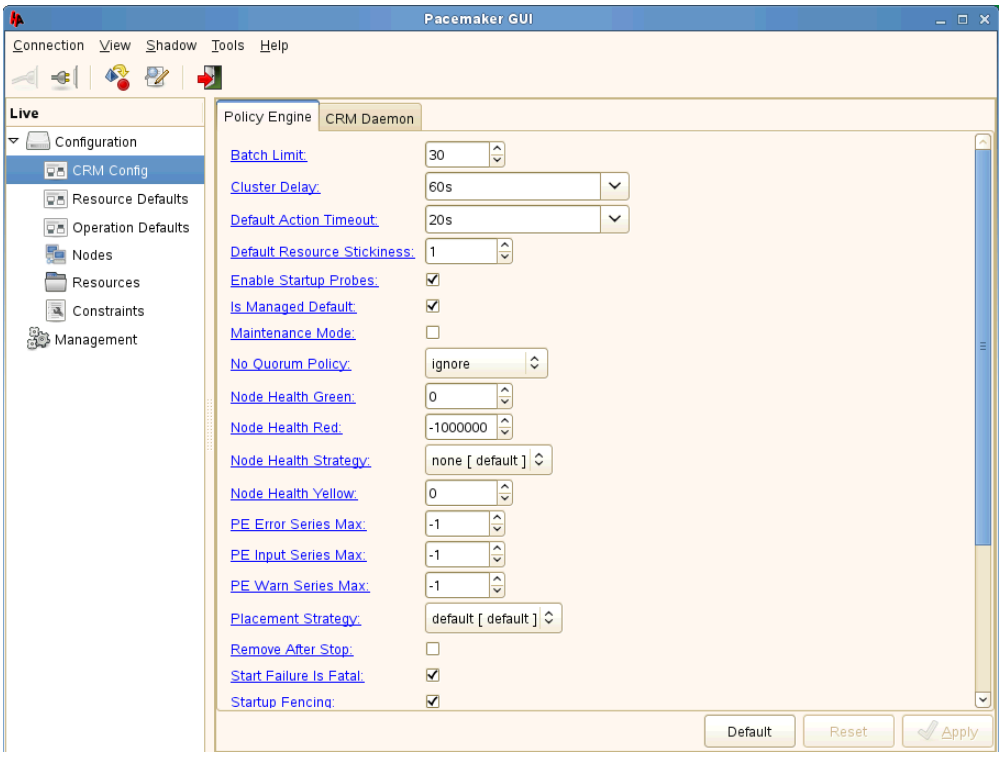

The value in this field controls how long Pacemaker waits for services to start. The default value is 20 seconds. The Orchestration Server requires more time than this to start. We recommend that you specify the value in this field at `120s`. More time might be required if your Orchestration Server grid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

1.2.4 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.

1.2.5 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 "Cloud Manager Orchestration Server Requirements" in the *NetIQ Cloud Manager Installation Planning Guide*.

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.

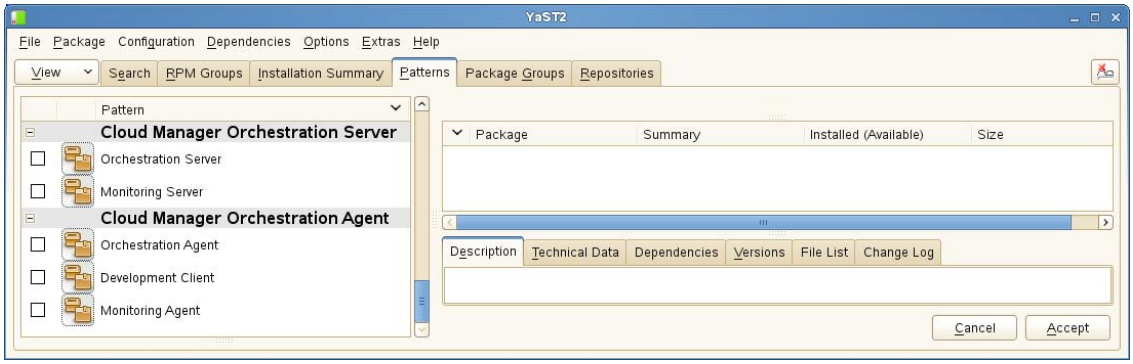

5. Select the Orchestration Server installation pattern for installation on the first node., then click *Accept*.

   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 19.

1.3 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).

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 1.4, "Installing and Configuring Orchestration Server Packages for High Availability on Other Nodes in the Cluster," on page 20.
2. Section 1.5, "Creating the Server Cluster Resource Group," on page 20.

3. Section 1.6, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 21
4. Section 1.7, “Installing and Configuring other Orchestration Components to the High Availability Grid,” on page 21.

1.3.1 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.

1.3.2 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	Your answer here determines whether this configuration takes place on a standard installation or on a High Availability installation. This section discusses standard installation, so specify `h` (for `ha` which means “high availability”).
Cluster Hostname or IP Address	Specify the fully qualified cluster hostname or the IP address that is used for configuring the Orchestration Server instance in a high availability cluster. The configuration script binds the IP address of the cluster to this server.

Server Configuration Requirement	Explanation and Action
Grid Name	A grid is an administrative domain container holding all of the objects in your network or data center. The Orchestration Server monitors and manages these objects, including users, resources, and jobs. The grid name you create here is displayed as the name for the container placed at the root of the Explorer tree in the Orchestration Console.
Administrator User	Specify a name for the Orchestration Server Administrator user. This name is used to log in as the administrator of the Orchestration Server and the objects it manages.
Administrator Password	Specify a password for the Orchestration Administrator user, then retype the password to validate it. You should remember this username for future logins.
Path to License File	A license key (90-day evaluation license or a full license) is required to use this product. You should have received this key from NetIQ, then you should have subsequently copied it to the network location that you specify here. Be sure to include the name of the license file in the path.
Auditing Database	We recommend that you do not install the audit database on this server.
Orchestration Agent Port[1]	Port 8100 is used for communication between the Orchestration Server and the Orchestration Agent. Specify another port number if 8100 is reserved for another use. If your Orchestration Server communicates with ESX servers, we recommend you configure port 8101. This requires that you configure all other Orchestration Agents communicating with this server to use port 8101. This configuration parameter is considered an advanced setting for the Orchestration Server in the GUI Configuration Wizard. If you select the *Configure Advanced Settings* check box in the wizard, you have the option of changing the default values. If you leave the check box deselected the setting is configured with normal defaults.
Administrator Information Port[1]	Port 8001 on the Orchestration Server provides access to an Administrator Information page that includes links to product documentation, agent and client installers, and product tools to help you understand and use the product. Specify another port number if 8001 is reserved for another use on this server.
TLS Certificate and Key[1]	Choose whether to generate a TLS certificate and key. ◆ Default = `yes` (the Orchestration Server must generate a certificate and key for authentication) ◆ A PEM-encoded TLS certificate and key is needed for secure communication between the Orchestration Server and Orchestration Agent. ◆ If you respond with `no`, you need to provide the location of an existing certificate and key.

Server Configuration Requirement	Explanation and Action
TLS Server Certificate[2]	Specify the full path to the TLS server certificate. ◆ Default = `/etc/ssl/servercerts/servercert.pem` ◆ Specify the path to the existing TLS certificate.
TLS Server Key[2]	Specify the full path to the TLS server private key. ◆ Default = `/etc/ssl/servercerts/serverkey.pem` ◆ Specify the path to the existing TLS private key.

[1] This configuration parameter is considered an advanced setting for the Orchestration Server in the Orchestration Components Configuration Wizard. If you select the *Configure advanced settings* check box in the wizard, the setting is configured with normal defaults. Leaving the check box deselected lets you have the option of changing the default value.

[2] This configuration parameter is considered an advanced setting for the Orchestration Server in the Orchestration Components Configuration Wizard. If you select the *Configure advanced settings* check box in the wizard, this parameter is listed, but default values are provided only if the previous value is manually set to no.

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

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

or

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

4 Continue with "Checking the Configuration" on page 19.

1.3.3 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 20.

1.3.4 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 1.2.5, "Installing the Orchestration Server on the First Clustered SLES 11 Node," on page 15

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.

1.4 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):

1.5 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)" 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 1.6, "Testing the Failover of the Orchestration Server in a High Availability Grid," on page 21).

1.6 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.

1.7 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 "Cloud Manager System Requirements" in the *NetIQ Cloud Manager 2.1.4 Application Server Installation Guide*.
- Use YaST2 to install the Orchestration packages of your choice to the network server resources of your choice. For more information, see "Installing Cloud Manager Orchestration Components" in the *NetIQ Cloud Manager 2.1.4 Orchestration Installation Guide*.

  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 "Installing Cloud Manager Orchestration Components" in the *NetIQ Cloud Manager 2.1.4 Orchestration Installation Guide*.

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).

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.

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 1.6, “Testing the Failover of the Orchestration Server in a High Availability Grid,” on page 21.

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

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.

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).

3 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.

3.1 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.

A 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.1 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).

A.2 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
`resource.vm.linuxha.cib_xml.id`	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.
`resource.vm.linuxha.cib_xml`	The stored CIB XML definition for this VM in the cluster. When a VM is discovered (that is, when the *Resync State*, *Check Status*, or *Discover VM Hosts and Repositories* actions are executed) this fact is overwritten with new data queried from the cluster CIB. For all other actions (for example, *Provision*, *Save Config*, etc.), the cluster's CIB definition for the VM is replaced with the contents of this fact.

- 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
`resource.vm.linuxha.cib_xml_default`	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 `resource.vm.linuxha.cib_xml` fact is empty).
`resource.vm.linuxha.clear_migrate_constraints_on_shutdown`	Used when shutting down (or restarting) the vm. If set, any location constraints created by migrate will be cleared at shutdown time. If there are specific location constraints created by an administrator, they are not cleared (unless they follow the naming convention for standard migrate constraints).
`resource.vm.linuxha.failcount.threshold`	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'.
`resource.vm.linuxha.migrate.timeout`	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 `vm.linuxha.vm_state.wait.timeout` (for migrate).
`resource.vm.linuxha.provision.timeout`	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 `vm.linuxha.vm_state.wait.timeout`, and is set as the start operation timeout.

Fact Name	Description
resource.vm.linuxha.shutdown.timeout	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 vm.linuxha.vm_state.wait.timeout, and is set as the stop operation timeout and the VM's shutdown_timeout in the CIB. To use per-VM values, set the value of this fact to the empty string (' '). The default is to use the value of the resource.vm.shutdown.timeout fact.
resource.vm.linuxha.vm_state.wait.interval	Defines the interval (in seconds) when the desired VM state should be re-checked (sleep interval). The default value is 5 seconds.
resource.vm.linuxha.vm_state.wait.timeout	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).

- 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
resource.linuxha.joblets.default_slots	Specifies the default number of joblets slots to create for a Xen Cluster resource.
resource.linuxha.cibadmin.cache.enabled	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.
resource.linuxha.cibadmin.cache.lifetime	The amount of time (in seconds) for which the cached CIB should stay valid.
resource.linuxha.cibadmin.cache.invalidate_on_write	Specifies (true or false) whether the CIB cache should be invalidated when writing to the CIB.

A.3 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-xenha2_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`).

A.4 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">
```