

# **SUSE Xen VM High Availability Configuration Guide**

**Cloud Manager 2.1.5**

January 31, 2013



## Legal Notice

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

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.

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

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.

Check Point, FireWall-1, VPN-1, Provider-1, and SiteManager-1 are trademarks or registered trademarks of Check Point Software Technologies Ltd.

Access Manager, ActiveAudit, ActiveView, Aegis, AppManager, Change Administrator, Change Guardian, Cloud Manager, Compliance Suite, the cube logo design, Directory and Resource Administrator, Directory Security Administrator, Domain Migration Administrator, Exchange Administrator, File Security Administrator, Group Policy Administrator, Group Policy Guardian, Group Policy Suite, IntelliPolicy, Knowledge Scripts, NetConnect, NetIQ, the NetIQ logo, PlateSpin, PlateSpin Recon, Privileged User Manager, PSAudit, PSDetect, PSPasswordManager, PSSecure, Secure Configuration Manager, Security Administration Suite, Security Manager, Server Consolidator, VigilEnt, and Vivinet are trademarks or registered trademarks of NetIQ Corporation or its affiliates in the USA. All other company and product names mentioned are used only for identification purposes and may be trademarks or registered trademarks of their respective companies.

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.

---

# Contents

<b>About This Guide</b>	<b>5</b>
<b>About NetIQ Corporation</b>	<b>7</b>
<b>1 Installing and Configuring the Orchestration Agent for Xen VM Deployment in a SLES HAE Cluster</b>	<b>9</b>
1.1 Xen Cluster Architecture .....	9
1.2 Installing the Orchestration Agent in a SLES 11 SP1 HAE Xen Cluster .....	10
1.3 Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster .....	11
1.3.1 Configuring the Agent for the Cluster .....	11
1.3.2 Creating the Agent Cluster Resource Group .....	13
1.3.3 Removing the Orchestration Agent from a Clustered VM Host .....	14
<b>A Sample Orchestration Agent CIB XML</b>	<b>15</b>



---

# About This Guide

This *High Availability Configuration Guide* provides the information for installing and configuring NetIQ Cloud Manager Orchestration Agent in a SLES 11 high availability environment. The guide provides information about the components and configuration steps necessary for preparing this environment, including instructions for configuring the Orchestration Agent in a cluster. The guide is organized as follows:

- ♦ [Chapter 1, “Installing and Configuring the Orchestration Agent for Xen VM Deployment in a SLES HAE Cluster,”](#) on page 9
- ♦ [Appendix A, “Sample Orchestration Agent CIB XML,”](#) on page 15

For information about installing and configuring the Orchestration Server in a high availability environment, see the *NetIQ Cloud Manager Orchestration Server High Availability Configuration Guide*.

## 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 2.1.5 documentation, see the [NetIQ Cloud Manager 2.x documentation site](https://www.netiq.com/documentation/cloudmanager2/) (<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](http://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

<b>Worldwide:</b>	<a href="http://www.netiq.com/about_netiq/officelocations.asp">www.netiq.com/about_netiq/officelocations.asp</a>
<b>United States and Canada:</b>	888-323-6768
<b>Email:</b>	<a href="mailto:info@netiq.com">info@netiq.com</a>
<b>Web Site:</b>	<a href="http://www.netiq.com">www.netiq.com</a>

## Contacting Technical Support

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

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

## 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](mailto: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 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.

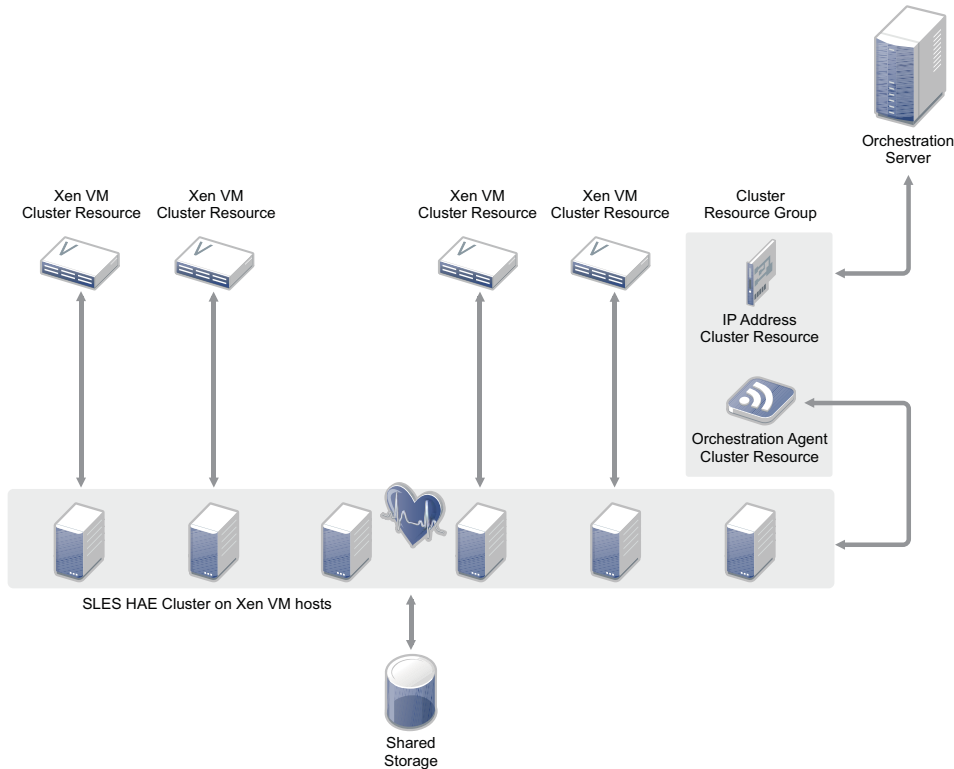
- ♦ [Section 1.1, “Xen Cluster Architecture,” on page 9](#)
- ♦ [Section 1.2, “Installing the Orchestration Agent in a SLES 11 SP1 HAE Xen Cluster,” on page 10](#)
- ♦ [Section 1.3, “Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster,” on page 11](#)

## 1.1 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 1-1** 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.

## 1.2 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 fills 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 1.3, “Configuring the Orchestration Agent in a SLES 11 SP2 HAE Xen Cluster,”](#) on page 11 to configure the Orchestration Agent.

## 1.3 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 1.3.1, “Configuring the Agent for the Cluster,”](#) on page 11
- ♦ [Section 1.3.2, “Creating the Agent Cluster Resource Group,”](#) on page 13
- ♦ [Section 1.3.3, “Removing the Orchestration Agent from a Clustered VM Host,”](#) on page 14

### 1.3.1 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: <a href="#">Agent Name</a> and <a href="#">Cluster Bind Address</a>.</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	<p>This value is required.</p>
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	<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>
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"> <li>◆ If an IP address was specified, the value is used unchanged.</li> <li>◆ 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.</li> <li>◆ 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.</li> </ul>
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><b>NOTE:</b> 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 1.3.2, “Creating the Agent Cluster Resource Group,” on page 13](#) 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.

## 1.3.2 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. A sample of this template is included in [Appendix A, “Sample Orchestration Agent CIB XML,”](#) on page 15.

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](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](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*.

### 1.3.3 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](https://www.suse.com/documentation/sle_ha/book_sleha/data/book_sleha.html)).

---

# A 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>
</group>
```

```

    </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>

```