

Orchestration Troubleshooting Reference

Cloud Manager 2.0

December 15, 2011



Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the [Novell International Trade Services Web page \(http://www.novell.com/info/exports/\)](http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2008-2011 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the [Novell Documentation Web page \(http://www.novell.com/documentation\)](http://www.novell.com/documentation/).

Novell Trademarks

For Novell trademarks, see [the Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Reference	5
1 Troubleshooting	7
1.1 Troubleshooting Installation Issues	7
1.2 Troubleshooting Upgrade Issues	8
1.3 Troubleshooting Orchestration Server Issues	8
1.4 Troubleshooting Orchestration Console Issues	9
1.5 Troubleshooting Orchestration Agent Issues	11
1.6 Troubleshooting VM Client Issues	11
1.7 Troubleshooting General VM Management Issues	14
1.8 Troubleshooting Xen VM Provisioning Actions	18
1.9 Troubleshooting vSphere VM Provisioning Actions	20
1.10 Troubleshooting Hyper-V VM Provisioning Operations	23
1.11 Troubleshooting Citrix Xen VM Provisioning Operations	26

About This Reference

This *Orchestration Troubleshooting Reference* provides troubleshooting information you need to identify and work around known issues in Orchestration components of NetIQ Cloud Manager.

- ♦ [Chapter 1, “Troubleshooting,” on page 7](#)

Audience

This book is intended for data center managers and IT or Operations administrators. The zos command line interface is intended for the use of Job Managers. The book assumes that users of the product have the following background:

- ♦ General understanding of network operating environments and systems architecture.
- ♦ Knowledge of basic UNIX shell commands and text editors.

Feedback

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, or go to www.novell.com/documentation/feedback.html (<http://www.novell.com/documentation/feedback.html>) and enter your comments there.

Documentation Updates

For the most recent version of this guide, visit the [NetIQ Cloud Manager 2.0 documentation Web site](http://www.novell.com/documentation/cloudmanager2/) (<http://www.novell.com/documentation/cloudmanager2/>).

Additional Documentation

For other NetIQ Cloud Manager documentation, visit the [NetIQ Cloud Manager 2.0 documentation Web site](http://www.novell.com/documentation/cloudmanager2/) (<http://www.novell.com/documentation/cloudmanager2/>).

1 Troubleshooting

This section provides information that identifies some ongoing known issues in NetIQ Cloud Manager Orchestration components and the methods that you can use to address those issues.

- ◆ [Section 1.1, “Troubleshooting Installation Issues,” on page 7](#)
- ◆ [Section 1.2, “Troubleshooting Upgrade Issues,” on page 8](#)
- ◆ [Section 1.3, “Troubleshooting Orchestration Server Issues,” on page 8](#)
- ◆ [Section 1.4, “Troubleshooting Orchestration Console Issues,” on page 9](#)
- ◆ [Section 1.5, “Troubleshooting Orchestration Agent Issues,” on page 11](#)
- ◆ [Section 1.6, “Troubleshooting VM Client Issues,” on page 11](#)
- ◆ [Section 1.7, “Troubleshooting General VM Management Issues,” on page 14](#)
- ◆ [Section 1.8, “Troubleshooting Xen VM Provisioning Actions,” on page 18](#)
- ◆ [Section 1.9, “Troubleshooting vSphere VM Provisioning Actions,” on page 20](#)
- ◆ [Section 1.10, “Troubleshooting Hyper-V VM Provisioning Operations,” on page 23](#)
- ◆ [Section 1.11, “Troubleshooting Citrix Xen VM Provisioning Operations,” on page 26](#)

1.1 Troubleshooting Installation Issues

The following sections provide solution to the problems you might encounter while performing the installation or configuration of the product:

- ◆ [“Configuration Programs Do Not Include a Way to Edit the Agent Configuration” on page 7](#)

Configuration Programs Do Not Include a Way to Edit the Agent Configuration

Source: Orchestration Installation and Configuration Programs.

Explanation: Although the scenario is not supported in a production environment, it is common in demonstration or evaluation situations to install the Orchestration Agent and the Orchestration Server on the same machine.

An error might occur if you install the agent after the initial server installation or if you attempt to use the configuration programs (`config`, `guiconfig`) to change the agent configuration after it is installed. Because of port checking routine in the configuration program, the error alerts you that port 8100 is already in use.

Action: To correct the problem for a demonstration setup, stop the Orchestration Server, configure the agent with one of the configuration programs, then restart the server.

1.2 Troubleshooting Upgrade Issues

The following sections provide solution to the problems you might encounter while performing an upgrade from an earlier version of the product:

- ♦ [“Currently Defined Job Schedule Deployment States Are Overwritten on Upgrade”](#) on page 8

Currently Defined Job Schedule Deployment States Are Overwritten on Upgrade

Source: Product upgrade process.

Explanation: The currently defined deployment state (that is, enabled or disabled) for a job schedule is overwritten by the default job deployment state when you upgrade from PlateSpin Orchestrate 2.6 to NetIQ Cloud Manager Orchestration Server 3.0.

Action: If you want to re-enable or disable a job after the upgrade, you need to open the Job Scheduler in the Orchestration Console and manually change the deployment state.

For more information, see [“The Orchestration Server Job Scheduler”](#) in the *NetIQ Cloud Manager 2.0 Orchestration Console Reference*.

1.3 Troubleshooting Orchestration Server Issues

The following sections provide solution to the problems you might encounter while using the Orchestration Server:

- ♦ [“Orchestration Server Might Appear to Be Deadlocked When Provisioning Large Numbers of Jobs with Subjobs”](#) on page 8
- ♦ [“Orchestration Server Might Hang if the System Clock Is Changed Abruptly”](#) on page 8
- ♦ [“Authentication to an Active Directory Server Might Fail”](#) on page 9
- ♦ [“The Orchestration Server Must Have Sufficient RAM”](#) on page 9

Orchestration Server Might Appear to Be Deadlocked When Provisioning Large Numbers of Jobs with Subjobs

Source: Cloud Manager Orchestration Server

Explanation: In some deployments where a large number of running jobs spawn subjobs, the running jobs might appear to stop, leaving jobs in the queue.

Possible Cause: This occurs because of job limits set in the Orchestration Server to avoid overload or “runaway” conditions.

Action: If this deadlock occurs, you can slowly adjust the job limits to tune them according to your deployment. For more information, see [“Job Limits Panel”](#) in the *NetIQ Cloud Manager 2.0 Orchestration Console Reference*.

Orchestration Server Might Hang if the System Clock Is Changed Abruptly

Source: Cloud Manager Orchestration Server

Explanation: As with many applications, you should avoid abrupt changes in the system clock on the machine where the Orchestration Server is installed; otherwise, the agent might appear to hang, waiting for the clock to catch up.

This issue is not affected by changes in clock time occurring from daylight saving adjustments.

Action: We recommend that you use proper clock synchronization tools such as a Network Time Protocol (NTP) server in your network to avoid large stepping of the system clock.

Authentication to an Active Directory Server Might Fail

Source: Cloud Manager Orchestration Server

Explanation: A simplified Active Directory Server (ADS) setup might be insufficient because of a customized ADS install (for example, `namingContexts` entries that generate referrals when they are looked up).

Possible Cause: The checking logic in the current AuthLDAP auth provider assumes that if any `namingContext` entry is returned, it has found the domain and it stops searching.

Action: If you encounter this issue, you need to manually configure LDAP as a generic LDAP server, which offers many more configuration options.

The Orchestration Server Must Have Sufficient RAM

Source: Cloud Manager Orchestration Server

Explanation: If the Orchestration Server fails to start after installation and configuration, sufficient RAM might not be installed on your hardware or assigned to the VM you are attempting to use.

Possible Cause: The Orchestration Server requires 3 GB of RAM to function with the preset defaults.

Action: If the server does not start, increase your physical RAM size (or, for a VM, increase the setting for virtual RAM size). Alternatively, you can reduce the JVM heap size, as explained in [“Validating and Optimizing the Orchestration Configuration”](#) in the *NetIQ Cloud Manager 2.0 Orchestration Installation Guide*.

1.4 Troubleshooting Orchestration Console Issues

The following sections provide solution to the problems you might encounter while using the Orchestration Console:

- ♦ [“The Server Console Displays Incorrect CPU Speed for SLES 11 SP1 Resources”](#) on page 9
- ♦ [“After Installing the Orchestration Agent on VM, the VM is Not Displayed as a Resource in the Orchestration Console”](#) on page 10

The Server Console Displays Incorrect CPU Speed for SLES 11 SP1 Resources

Source: Orchestration Console

Explanation: The CPU speed displayed in the Orchestration Console (see the `resource.cpu.mhz` and `resource.metrics.cpu_speed` facts) for SLES 11 SP1 resources is incorrect. The invalid display results from powersave settings on the CPU. Until the CPU has been run at full speed, `/proc/cpuinfo` displays this incorrect value for CPU MHz, and the value in the Orchestration Server is also incorrect.

Possible Cause: The issue results from the CPU starting in powersave mode. This slows down the CPU until it is needed, so `/proc/cpuinfo` does not show the maximum potential speed of the CPU. Instead, it shows the maximum speed that the CPU has shown since boot time.

Action: To work around this issue, run the `powersave --performance-speed` command at the server command line.

This command forces the CPU to reach its maximum speed, so you should see the correct value displayed in `/proc/cpuinfo` and the Development Client should also display the correct speed. After you run this command, you can set the powersave mode to a normal state with either of the following commands:

```
powersave --powersave-speed
```

or

```
powersave --dynamic-speed
```

When the powersave mode is set to a normal state, `/proc/cpuinfo` retains the accurate value for the current CPU speed.

TIP: To see the contents of `/proc/cpuinfo`, run the `cat /proc/cpuinfo` command at the bash prompt of your SLES server.

After Installing the Orchestration Agent on VM, the VM is Not Displayed as a Resource in the Orchestration Console

Source: The Orchestration Console

Action: Do the following:

- ◆ Ensure that the Orchestration Agent is running on the VM.
- ◆ Ensure that no errors have been logged into the `agent.log` file.

The log file is located in the

`<Orchestration_Agent_installation_directory>\novell\zos\agent\node.default` directory on Windows and in the `/opt/novell/zos/agent/node.default` directory on Linux.

- ◆ Ensure that the Orchestration Server is registered to the DNS server.

1.5 Troubleshooting Orchestration Agent Issues

The following sections provide solution to the problems you might encounter while using the Orchestration Console:

- ♦ [“Orchestration Agent Fails to Set the UID on Files Copied from the Datagrid” on page 11](#)

Orchestration Agent Fails to Set the UID on Files Copied from the Datagrid

Source: Orchestration Agent

Explanation: If Network File System (NFS) is used to mount a shared volume across nodes that are running the Orchestration Agent, the agent cannot properly set the UID on files copied from the datagrid to the managed nodes by using the default NFS configuration on most systems.

Action: To address this problem, disable root squashing in NFS so that the agent has the necessary privileges to change the owner of the files it copies.

For example, on a Red Hat Enterprise Linux (RHEL) NFS server or on a SUSE Linux Enterprise Server (SLES) NFS server, the NFS configuration is set in `/etc/exports`. The following configuration is needed to disable root squashing:

```
/auto/home * (rw, sync, no_root_squash)
```

In this example, `/auto/home` is the NFS mounted directory to be shared.

NOTE: The GID is not set for files copied from the datagrid to an NFS mounted volume, whether root squashing is disabled or not. This is a limitation of NFS.

1.6 Troubleshooting VM Client Issues

The following sections provide solution to the problems you might encounter while using the Cloud Manager VM Client:

- ♦ [“VM does not start when a local repository is assigned to multiple hosts” on page 11](#)
- ♦ [“Not configuring a display driver triggers a pop-Up message” on page 12](#)
- ♦ [“Cannot increase the number of vCPUs on a running Xen VM” on page 12](#)
- ♦ [“The default desktop theme on SLES 10 or SLED 10 causes a display problem for the VM Client” on page 12](#)
- ♦ [“Using the Cloud Manager VM Client in a firewall environment” on page 13](#)
- ♦ [“VM Client error log lists a high number of display exceptions” on page 13](#)
- ♦ [“Storage type options might not be visible when modifying a repository” on page 13](#)
- ♦ [“The Orchestration Console must be used to set the administrator and domain facts before cloning” on page 13](#)
- ♦ [“Xen VMs created in the VM Client are not created in the correct location” on page 14](#)

VM does not start when a local repository is assigned to multiple hosts

Source: Cloud Manager VM Client

Explanation: When you configure local repositories in the VM Client, the program does not check to verify that it is set up correctly on the server.

Make sure that if you associate a repository to a host that it actually has access and rights to use that repository. Otherwise, if a VM attempts to start on a host without access to the repository, it does not start and no longer displays in the VM Client or Orchestration Console. You can recover from this situation by fixing the repository access and rediscovering the VMs.

An example of this would be a Linux host that is associated to a NAS repository but has not been granted access to the NFS server's shared directory.

Action: To work around this issue, correctly set up your local repositories on your host servers, and do not share the local repositories. Allow only the host server that owns the local repository to have access to it.

Not configuring a display driver triggers a pop-Up message

Source: Cloud Manager VM Client

Explanation: If you configure a VM with `None` for the display driver and select to install the VM, a VNC pop-up window displays, but the VNC is never connected.

Action: To work around this issue, be careful not to configure a VM without a display driver. You can also connect to the VM using `ssh` or some other utility.

Cannot increase the number of vCPUs on a running Xen VM

Source: Cloud Manager VM Client

Explanation: The vCPUs number that you set on a Xen VM is the maximum number of vCPUs allowed for that instance of the VM when you run it.

The VM Client allows you to increase the number of vCPUs beyond the originally defined number while a VM is running. However, these "extra" vCPUs (the number of vCPUs over the initial amount) are not recognized by Xen.

Therefore, when using `Apply Config` to modify the number of vCPUs on a running VM instance, the number can be less than or equal to, but not greater than the initial number set when the VM instance was started.

Action: To work around this issue do not use `Apply Config` to increase the number of vCPUs higher than the originally defined number for the Xen VM instance when it was provisioned.

The default desktop theme on SLES 10 or SLED 10 causes a display problem for the VM Client

Source: Cloud Manager VM Client

Explanation: If you edit the details for a storage (repository) item in the VM Client, such as changing the path, nothing appears in the combo box (you see only white space). The display problem is caused by a conflict with the default desktop theme installed with SLES 10 or SLED 10.

Action: You can work around this issue by changing the SLES 10 or SLED 10 desktop theme:

- 1 On the SLES or SLED desktop, click the *Computer* icon on the lower left to open the Applications dialog box.
- 2 In the Applications dialog box, click *More Applications* to open the Applications Browser.
- 3 In the left panel of the Applications Browser, click *Tools* to go to the *Tools* menu in the browser.
- 4 In the *Tools* menu, select *Control Center* to open the Desktop Preferences dialog box.
- 5 In the *Look and Feel* section of the preferences menu, select *Theme* to open the Theme Preferences dialog box.
- 6 Select any theme other than the current SLES or SLED default, then click *Close*.

Using the Cloud Manager VM Client in a firewall environment

Source: Cloud Manager VM Client

Explanation: Using the Cloud Manager VM Client in a NAT-enabled firewall environment is not supported for this release. The Orchestration Console uses RMI to communicate with the server, and RMI connects to the initiator on dynamically chosen port numbers. To use the VM Client in a NAT-enabled firewall environment, you need to use a remote desktop or VPN product.

Action: If you are using a firewall that is not NAT-enabled, the VM Client can log in through the firewall by using port 37002.

VM Client error log lists a high number of display exceptions

Source: Cloud Manager VM Client

Explanation: A large number of exceptions involving the `org.eclipse.ui` plug-in are listed in the VM Client error log. These errors originate from some of the Eclipse libraries used by the VM Client.

Action: NetIQ is aware of the high number of exceptions occurring within this class. The errors are currently unavoidable and can be safely ignored.

Storage type options might not be visible when modifying a repository

Source: Cloud Manager VM Client

Explanation: While you are modifying a Storage Repository in the VM Client interface on a Linux desktop, you might have difficulty seeing different storage type options because of a font color in the display. The problem is not seen on all machines where the VM Client can be installed.

The Orchestration Console must be used to set the administrator and domain facts before cloning

Source: Cloud Manager VM Client

Explanation: The *Network* and *Windows* tabs have been removed from the Clone VM Wizard in the VM Client. You need to use the Orchestration Console to set the Administrator and Domain facts prior to cloning in the VM Client. In addition, because the *Windows* tab no longer exists, the *Use Autoprep* option is always set to True when cloning from the VM Client.

Xen VMs created in the VM Client are not created in the correct location

Source: Cloud Manager VM Client

Explanation: When you use the VM Client to create a new Xen VM, the client places that VM at the root of the file system instead of the selected repository path.

For example, if the repository path is `/var/lib/xen/images/` and the specified path for disk images in the VM Client is `testvm`, the client puts the VM in the `/testvm` directory rather than in the `/var/lib/xen/images/testvm` directory as you would intuitively expect.

Action: To work around this issue, when you use the VM Client to build a VM, specify the absolute path to the desired location for the disk image

1.7 Troubleshooting General VM Management Issues

The following sections provide solutions to the problems you might encounter while working with general VM management operations:

- ♦ [“Volume Tools Hang While Scanning a Suspended Device” on page 14](#)
- ♦ [“SUSE Linux VMs Might Attempt To Partition a Read-only Device” on page 15](#)
- ♦ [“RHEL 5 VMs Running the Kudzu Service Do Not Retain Network Interface Changes” on page 15](#)
- ♦ [“Policies Applied to VM Resources Are Deleted” on page 15](#)
- ♦ [“VMs Provisioned from a VM Template Are Not Restarted When a VM Host Crashes” on page 17](#)
- ♦ [“Unable to Provision a VM to Another Cluster Node Due to Reason “VM Networks Are Not Available”” on page 17](#)

Volume Tools Hang While Scanning a Suspended Device

Source: Scanned device.

Explanation: When a mapped device is in a suspended state, volume tools such as `vgscan`, `lvscan`, and `pvscan` hang. If the `vmprep` job is run on such a device, it throws an error such as the following to alert you to the condition:

```
vmquery: /var/adm/mount/vmprep.df8fd49401e44b64867f1d83767f62f5:
Failed to
mount vm image "/mnt/nfs_share/vms/rhel4tmpl2/disk0": Mapped
device
/dev/mapper/loop7p2 appears to be suspended. This might cause
scanning for
volume groups (e.g. vgscan) to hang.
WARNING! You may need to manually resume or remove this mapped
device (e.g.
dmsetup remove /dev/mapper/loop7p2)!
```

Action: Because of this behavior, we recommend against using LVM and similar volume tools on a virtual machine managed by Orchestration Services.

SUSE Linux VMs Might Attempt To Partition a Read-only Device

Source: YaST Partitioner.

Explanation: When you build a SUSE Linux VM and specify a read-only virtual device for that VM, in some instances the YaST partitioner might propose a re-partitioning of the read-only virtual device.

Possible Cause: Although Xen normally attempts to notify the guest OS kernel about the mode (`ro` or `rw`) of the virtual device, under certain circumstances the YaST partitioner proposes a re-partitioning of the virtual device that has the most available disk space without considering the other device attributes. For example, if a specified CD-ROM device happens to be larger than the specified hard disk device, YaST attempts to partition the CD-ROM device, which causes the VM installation to fail.

Action: To work around this issue, connect a VNC console to the VM being built during the first stage of the VM install, then verify the partition proposal before you continue with the installation. If the partition proposal has selected an incorrect device, manually change the selected device before you continue with the installation of the VM.

RHEL 5 VMs Running the Kudzu Service Do Not Retain Network Interface Changes

Source: Kudzu service.

Explanation: Anytime you modify the hardware configuration (for example, changing the MAC address or adding a network interface card) of a RHEL 5 VM that is running the Kudzu hardware probing library, the VM does not retain the existing network interface configuration.

Possible Cause: When you start a RHEL 5 VM, the Kudzu service recognizes the hardware changes at boot time and moves the existing configuration for that network interface to a backup file. The service then rewrites the network interface configuration to use DHCP instead.

Action: To work around this problem, disable the Kudzu service within the RHEL VM by using the `chkconfig --del kudzu` command.

Policies Applied to VM Resources Are Deleted

Source: VM clones awaiting provisioning.

Explanation: Provisioning code requires that VMs and VM clones be standalone (that is, they are [removed from a template dependency](#) and are no longer considered to be “linked clones”).

Possible Cause: VMs in PlateSpin Orchestrate 2.5 and later must be made standalone to receive and retain associated policies.

Action: Apply a conditional policy to the parent template that can be applied to the clones while they are running. Depending upon the facts set on the clone, the inherited VM host constraint can be conditionally applied to the clone.

The following is an example of a conditional policy that you could apply to the VM template to restrict vmhost based on resource attributes (group membership, etc.).

```
<policy>
  <constraint type="vmhost">
    <if>
      <contains fact="resource.groups" value="exclude_me"
        reason="Only apply this vmhost constraint to
resources NOT in exclude_me resource group" >
      </contains>
    <else>
      <if>
        <defined fact="resource.some_boolean_fact" />
        <eq fact="some_boolean_fact" value="true" />
        <then>
          <contains fact="vmhost.resource.groups"
value="first_vmhost_group"
          reason="When a resource is not in
the exclude_me group, when some_boolean_fact is true,
          provision to a vmhost in
the first_vmhost_group"/>
        </then>
        <else>
          <if>
            <defined
fact="resource.some_other_boolean_fact" />
            <eq fact="some_other_boolean_fact"
value="true" />
            <not>
              <and>
                <eq fact="resource.id"
value="never_use_this_resource"
                reason="Specifically exclude
this resource from consideration." />
              <or>
                <eq fact="vmhost.cluster"
factvalue="resource.provision.vmhost.cluster" />
                <eq fact="vmhost.cluster"
factvalue="resource.provision.vmhost" />
              </or>
            </and>
          </not>
          <then>
            <contains
fact="vmhost.resource.groups" value="another_vmhost_group"
            reason="When a resource is
not in the exclude_me group, when some_boolean_fact is false,
            and
some_other_boolean_fact is true, (but also not some other things),
            provision to a vmhost
in another_vmhost_group"/>
          </then>
        </if>
      </else>
    </if>
  </else>
</if>
</constraint>
</policy>
```

VMs Provisioned from a VM Template Are Not Restarted When a VM Host Crashes

Source: VM host with VMs provisioned from a template.

Explanation: If a VM host crashes, VMs that were provisioned from a template on that host are not restarted on another active VM host. Instead, the Orchestration Server provisions another VM cloned from the original template, on the next available host. The disk files of the original clone are not destroyed (that is, “cleaned up”) after the crash, but the original VM template files are destroyed.

If a Discover Repository action is issued before the cloned VM is deleted from the crashed host, the Orchestration Server creates a new VM object with the `zombie_` string prepended to the VM object name.

Possible Cause: While hosting a provisioned clone, VM host crashed or the Orchestration Agent on that host went offline.

Action: To work around this issue, you can either remove the VM from the file system before the Orchestration Server rediscovers it, or you can issue a Destroy action on the discovered “zombie” VM.

Unable to Provision a VM to Another Cluster Node Due to Reason “VM Networks Are Not Available”

Source: The Orchestration Console

Explanation: A prerequisite for clustering is that every node contained within a cluster should be symmetric. That is, every node in a cluster should have visibility to all the networks and storage provided by the cluster. In this case, because a VM host cluster must be able to place the VM on any node in the cluster, the networks shown as being available to that cluster are the intersection of all the networks available on the VM host nodes that are members of the cluster (see the `vmhost.networks` fact on the cluster object).

Action: Reconfigure each of the cluster nodes to provide the networks required by the VM host cluster and re-run the *Discover Hosts* action.

Alternatively, you can reconfigure the VM to use another network available to all cluster nodes. After you choose a new network(s) configuration for a VM, make sure you run the *Save Config* action to commit these changes to the VM configuration.

When you reconfigure the networks on a VM, at least one network option, *all*, is available. This option designates that any network can be suitable for VM placement. Choosing this option allows the network constraint to pass, and the provisioning adapter is then responsible for configuring a new network as it sees fit.

1.8 Troubleshooting Xen VM Provisioning Actions

The following sections provide solution to the problems you might encounter while performing provisioning actions on VMs created in SUSE Xen and managed by the Orchestration Server:

- ♦ [“Provisioning a Xen VM Does Not Work on the Host Server” on page 18](#)
- ♦ [“Multiple Instances of the Same Xen VM Running when Located on Shared Storage” on page 18](#)
- ♦ [“Running xm Commands on an Old Xen VM Host Causes Server to Hang” on page 19](#)

Provisioning a Xen VM Does Not Work on the Host Server

Source: The Orchestration Console

Explanation: When you try to provision a Xen VM, the job might fail with the following error message in the job log:

```
[c121] RuntimeError: vmprep: Autoprep of /var/lib/xen/images/min-  
tmpl-1-2/disk0  
failed with return code 1: vmprep: autoprep:  
/var/adm/mount/vmprep.3f96f60206a2439386d1d80436262d5e: Failed to  
mount vm  
image "/var/lib/xen/images/min-tmpl-1-2/disk0": vmmount: No root  
device found  
Job 'zosSystem.vmprep.76' terminated because of failure. Reason:  
Job failed
```

A VM host cannot provision a VM that has a different file system than the VM host. The currently supported file systems are ext2, ext3, reiserfs, jfs, xfs, vfat, and ntfs.

Action: To work around the issue, load the VM’s file system Linux module on the VM host, or add it to the Linux kernel if a custom kernel is being used.

Typically, Linux kernels autoload the appropriate module to do the work.

You must manually load the proper kernel module on the VM host to support the VM’s file system.

For example, if the VM host uses ext3 and the VM image uses reiserfs, load the proper kernel module onto the VM host to support the VM image’s reiserfs file system. Then, on the VM host, run:

```
modprobe reiserfs
```

Next, provision the VM.

NOTE: Cloning with prep is limited to what the Virtual Center of VMware Server supports.

Multiple Instances of the Same Xen VM Running when Located on Shared Storage

Source: Shared storage for Xen VMs.

Explanation: The `xendConfig` job runs when a VM host is added to the Orchestration Server. This job automates some of the configurations possible on a Xen VM host. With the default Xen configuration, it is possible to incorrectly start a running VM a second time from storage that is shared by and accessible to another Xen VM host.

Possible Cause: A running Xen VM can only be locked to a specific Xen VM host when the `xend` service is configured to share a VM domain lock file on a shared file system. By default, the `xend` service places these VM domain lock files in the `/var/lib/xend/domains` directory, which is usually not located on shared storage.

Action: You can configure Xen VM locks in the Orchestration Server by uncommenting certain facts in the policy file (search for `xend.xend-domain-lock`).

Uncomment these facts in `xendConfig.policy`:

```
<!--
  <fact name="xend.xend-domain-lock"
        type="String"
        description="Create an external lock file when domains are started. Lock
                    file is placed in xend-domains-lock-path/[domain_uuid] on domain
                    startup and removed when domain is stopped. Set to yes to enable
                    lock file creation.

                    Note that external locking mechanisms are no substitute for a cluster
                    environment that protects shared resources, but may be useful in
                    some circumstances nonetheless."
        value="yes" />
  <fact name="xend.xend-domain-lock-path"
        type="String"
        description="Path where domain lock is stored if xend-domain-lock is enabled.
                    Note: This path must be accessible to all VM Servers participating
                    in domain locking, e.g. by specifying a shared mount point.
                    Lock is placed in [xend-domain-lock-path]/[domain-uuid].
                    Default is /var/lib/xend/domains/"
        value="/var/lib/xend/domains" />
-->
```

To uncomment a section of code, remove the “<!--” (comment open) tag and the “-->” (comment close) tag. Edit the `xend-domain-lock-path` fact to set an alternate location on shared storage that is available to all VM hosts.

When you make the changes and save the file, the facts become active and the VM locking parameters of each newly joining VM host are adjusted accordingly.

You can also schedule an immediate run of the `xendConfig` job to adjust all configuration files of the Xen VM hosts that are already connected to the Orchestration Server.

NOTE: Setting the lock path by using the Orchestration Server only supports the scenario where all VM hosts have the domain lock path connected to the same shared repository. For more complex setups, you need to use alternative methods to adjust the VM host lock configurations.

Running `xm` Commands on an Old Xen VM Host Causes Server to Hang

Source: The source of the message.

Explanation: The Xen provisioning adapter uses `xm` commands to perform basic VM life cycle operations such as building a VM, starting a VM, stopping a VM, pausing a VM, and suspending a VM. These commands can cause the server to hang if it has not been updated with the latest Xen tools.

Action: Make sure the Xen VM host has the latest Xen tools available by running the `rpm -qa | grep xen-tools` command.

You should have the SLES 11 Xen maintenance release #1 (or later) of the tools:

```
Xen 3.3.1_18546_14
```

1.9 Troubleshooting vSphere VM Provisioning Actions

The following sections provide solution to the problems you might encounter while performing provisioning actions on VMs managed by the VMware vCenter hypervisor:

- ♦ [“Unable to perform any Provisioning Adapter Action after the Save Config Action on the vSphere Managed VM” on page 20](#)
- ♦ [“\(503\) Service Unavailable Errors Might Occur while Cloning vSphere VMs” on page 20](#)
- ♦ [“Invalid Datastore Path Error” on page 21](#)
- ♦ [“Running Provisioning Operations on a Batch of vSphere VMs Results in JDL Event Handler Errors” on page 21](#)

Unable to perform any Provisioning Adapter Action after the Save Config Action on the vSphere Managed VM

Source: The Orchestration Console.

Possible Cause: The VM UUID value of the vSphere managed VM is not a 128-bit hexadecimal value. Even though the Save Config action is successful and the VM is provisioned, the hypervisor automatically assigns a different UUID value. Subsequently, any provisioning adapter action performed on the VM fails.

Action: Specify a 128-bit hexadecimal value for the VM UUID.

- 1 In the Orchestration Console, click *Resources* > the vSphere managed VM. The *Info/Groups* tab is displayed by default.
- 2 In the Virtual Machine Configuration panel, set the value of *VM UUID* to a 128-bit hexadecimal value.
- 3 Right-click the vSphere managed VM, then click *Save Config*.

(503) Service Unavailable Errors Might Occur while Cloning vSphere VMs

Source: The Orchestration Console

Explanation: Running the *Clone* action repeatedly on vSphere VM templates might result in the following error:

```
Clone : (503)Service Unavailable
```

Possible Cause: This error indicates that the server is currently unable to handle the request due to a temporary overloading or maintenance of the server. Testing has shown that this error is most likely to occur when vSphere and the Orchestration Agent are both installed on the same Windows Server 2003 computer.

Action: If you encounter this error, we recommend that you download and apply the appropriate [Microsoft hotfix \(http://support.microsoft.com/kb/979230\)](http://support.microsoft.com/kb/979230) to the vCenter server.

Invalid Datastore Path Error

Source: The Orchestration Server.

Explanation: When attempting to *Save Config* a vSphere VM with an ISO-backed vDisk (for example, a vDisk that specifies a location in the /vmimages folder and does not have its repository fact set), the job fails with a message similar to the following:

```
VMSaveConfig : Invalid datastore path '/vmimages/tools-isoimages/linux.iso'
```

Action: To work around this issue, associate a policy with the ISO-backed vdisk object that prepends an empty datastore string ([]) to the beginning of the `vdisk.location` fact. For example:

```
<policy>
  <vdisk>
    <fact name="location"
          type="String"
          value="[] /vmimages/tools-isoimages/linux.iso" />
  </vdisk>
</policy>
```

Running Provisioning Operations on a Batch of vSphere VMs Results in JDL Event Handler Errors

Source: The Orchestration Console

Explanation: If you write JDL scripts to automate provisioning actions for a large number of vSphere VMs, you might receive a failure notice similar to the following:

```
Job 'testadmin.r_testvm_resync_batch.15684' terminated because of failure.
Reason: job exceeded max limit of jdl event handler
Job 'testadmin.r_testvm_resync_batch.15684' terminated because of failure.
Reason: job exceeded max limit of jdl event handler
```

You also see the following error in `server.log`:

```
08.24 17:32:59: JobManager,NOTICE: job instance 'testadmin.r_testvm_resync_batch.15082' failed
08.24 17:46:25: JobManager,NOTICE: job instance 'testadmin.r_testvm_resync_batch.15684' failed
08.24 17:46:25: Broker,ERROR: Exception in thread "JDL Event (job failed event) jobId (testadmin.r_testvm_resync_batch.15684)"
08.24 17:46:25: Broker,ERROR: ValueError: I/O operation on closed file
```

Possible Cause: This error indicates that maximum number of JDL threads allowed by the server have been exceeded. Testing has shown that numerous instances of the `provisioner_completed_event` are blocked and waiting for the `provisioner_job_to_finish_its_job_started_event`.

Action: Rewrite the original script. The original script might look like this:

```

import time
class testvm_resync(Job):
    def job_started_event(self):
        vms_group = getMatrix().getGroup(TYPE_RESOURCE, 'VMs') #
gets the matrix object id for 'VMs' group
        vms = vms_group.getMembers() # gets the group members of
'VMs' group
        for vm in vms:
            id = vm.getFact("resource.id") #gets the resource.id
fact of a vm
            thevmtype = vm.getFact("resource.type") # find the vm
type
            if id.startswith("c-") and thevmtype == 'VM': # search
criteria
                vmstate = vm.getFact("resource.provision.state") #
find the vm state
                thevm = getMatrix().getGridObject(TYPE_RESOURCE,
id); #gets the vm's id
                thevm.check() # vm life cycle operations
                time.sleep(2*60) #pause time - 2 min

```

The rewritten script might look like this:

```

import time
class testvm_resync(Job):
    def job_started_event(self):
        timer = Timer(self.prov,0)
    def prov(self):
        vms_group = getMatrix().getGroup(TYPE_RESOURCE, 'VMs') #
gets the matrix object id for 'VMs' group
        vms = vms_group.getMembers() # gets the group members of
'VMs' group
        for vm in vms:
            id = vm.getFact("resource.id") #gets the resource.id
fact of a vm
            thevmtype = vm.getFact("resource.type") # find the vm
type
            if id.startswith("c-") and thevmtype == 'VM': # search
criteria
                vmstate = vm.getFact("resource.provision.state") #
find the vm state
                thevm = getMatrix().getGridObject(TYPE_RESOURCE,
id); #gets the vm's id
                thevm.check() # vm life cycle operations
                time.sleep(2*60) #pause time - 2 min

```

This change lets the `job_started_event` end after transferring the process to another JDL event/method to run on a timer basis. In this example, the timer is set for 10 seconds, but you could set it to zero.

The timer is normally used for callback. For example, the vSphere provisioning adapter uses `Timer` to check every 30 seconds whether a vSphere action is still working or dead.

This not isolated to the `check()` action. It includes other actions such as `provision()`, `shutdown()`, `suspend()`, `checkpoint()`, `saveConfig()` and `restart()`.

1.10 Troubleshooting Hyper-V VM Provisioning Operations

The following sections provide solution to the problems you might encounter while performing provisioning operations on VMs managed by the Hyper-V hypervisor:

- ♦ [“The VM is Suspended When you Try to Revert the Snapshot of a Powered-on VM Running on a Hyper-V host” on page 23](#)
- ♦ [“Hyper-V Provisioning Jobs Fail When Several Jobs Are Started Simultaneously” on page 23](#)
- ♦ [“Hyper-V Provisioning Adapter Does Not Discover iSCSI Target Repositories” on page 24](#)
- ♦ [“Creating a VM with a Unique Path Results in VM Not Discovered” on page 24](#)
- ♦ [“OS info for Hyper-V VMs is Not Always Auto-Discovered by the Orchestration Agent.” on page 25](#)
- ♦ [“After Adding a vDisk to a Hyper-V Linux VM, That vDisk Is Not Visible Inside the VM” on page 25](#)
- ♦ [“Sysprep Does Not Work” on page 25](#)
- ♦ [“VNC Console for Hyper-V VM Does Not Open” on page 25](#)

The VM is Suspended When you Try to Revert the Snapshot of a Powered-on VM Running on a Hyper-V host

Source: Hyper-V provisioning adapter job

Explanation: If you try to revert the snapshot of a powered-on VM running on a Hyper-V host, the VM is suspended. This is a known behavior of VMs running on a Hyper-V host.

Action: Provision the suspended VM:

- 1 In the Orchestration Console, right-click the suspended VM, then click *Provision*.

The Provision VM dialog box is displayed.

- 2 In the *Plan (Host/Repository)* drop-down list, select the appropriate Hyper-V host.
- 3 Click *OK*.

Hyper-V Provisioning Jobs Fail When Several Jobs Are Started Simultaneously

Source: Hyper-V provisioning adapter job

Explanation: If you start more than the default number of Hyper-V provisioning jobs at the same time (for example, creating a template on each of three Hyper-V VMs simultaneously), the jobs fail because of an insufficient number of joblet slots set aside for multiple jobs.

Action: If you need to run more than the default number of joblets (one is the default for Hyper-V) at one time, change the *Joblet Slots* value on the VM host configuration page, or change the value of the `joblet.maxwaittime` fact in the hyperv policy so that the Orchestration Server waits longer to schedule a joblet before failing it on the VM host because of no free slots.

For more information, see “[Joblet Slots](#)” in the “[The Resource Object](#)” section of the *NetIQ Cloud Manager 2.0 Orchestration Console Reference*.

Hyper-V Provisioning Adapter Does Not Discover iSCSI Target Repositories

Source: Hyper-V provisioning adapter job

Explanation: If you are managing Windows VMs in a Hyper-V environment (clustered or non-clustered), the Cloud Manager Orchestration Server hyperv provisioning adapter fails 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`.

Action: To work around the issue:

- 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.

Creating a VM with a Unique Path Results in VM Not Discovered

Source: Orchestration Server

Explanation: 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.

Action: To work around this issue, you can edit the preferred path for the discovered repository where the VM resides, or you can create a new repository in the Orchestration Console with the preferred path to the Hyper-V VM.

OS info for Hyper-V VMs is Not Always Auto-Discovered by the Orchestration Agent.

Source: The Orchestration Agent

Explanation: As with other VMs managed by Xen and VMware, the OS info for Hyper-V VMs is not always auto-discovered by the Orchestration Agent.

The OS info for Hyper-V VMs is discovered only in the following circumstances:

- ♦ The Discover VM Images event triggers a Resync event for offline VMs to get OS info. However, OS info is not retrieved for the discovered templates.
- ♦ The Resync event on the VM retrieves the OS family and type if it is offline.
- ♦ The Resync event on the template sets the OS family to “Windows” because only Windows templates are supported in the hyperv provisioning adapter.
- ♦ The Create Template event tries to retrieve OS info before creating the template. Create Template succeeds only if the VM’s OS family is Windows.

Action: To work around this issue, you can either enter this information manually, install the Orchestration Agent on the VM to enable discovery, or use the Key/Value Pair Exchange mechanism to support integration services installation.

After Adding a vDisk to a Hyper-V Linux VM, That vDisk Is Not Visible Inside the VM

Source: The Orchestration Console

Explanation: After adding a vDisk to a Hyper-V Linux VM, that vDisk is not visible inside the VM.

Action: To work around this issue, 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.

Sysprep Does Not Work

Source: The Orchestration Console

Explanation: Sysprep does not work on Hyper-V Windows VMs.

Action: 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](#)” in the *NetIQ Cloud Manager 2.0 Orchestration Installation Guide*.

VNC Console for Hyper-V VM Does Not Open

Source: The Orchestration Console

Explanation: 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.

Action: 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*.

1.11 Troubleshooting Citrix Xen VM Provisioning Operations

The following sections provide solution to the problems you might encounter while performing provisioning operations on VMs managed by the Citrix Xen hypervisor:

- ♦ [“The Move, Create Template, and Clone Operations on a Citrix Xen VM Fail Occasionally” on page 26](#)

The Move, Create Template, and Clone Operations on a Citrix Xen VM Fail Occasionally

Source: Orchestration Server.

Explanation: The following error message is occasionally displayed when you perform a Move, Create Template or Clone operation on a Citrix Xen VM using the Orchestration Console:

```
The server failed to handle your request, due to an internal error. The given message may give details useful for debugging the problem.  
[<hostname>] message: Xmlrpcclient.Http_header_truncated("")
```

Action: Reboot the VM host machine to reset the condition.