

Cloud Manager Troubleshooting Reference

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

About This Reference	5
About NetIQ Corporation	7
1 Troubleshooting Cloud Manager Orchestration Issues	9
1.1 Troubleshooting Installation Issues	9
1.2 Troubleshooting Orchestration Server Issues	10
1.3 Troubleshooting Orchestration Console Issues	12
1.4 Troubleshooting Orchestration Agent Issues	13
1.5 Troubleshooting General VM Management Issues	13
1.6 Troubleshooting vSphere VM Provisioning Actions	18
1.7 Troubleshooting Citrix Xen VM Provisioning Operations	21
1.8 Troubleshooting Hyper-V VM Provisioning Operations	21
1.9 Troubleshooting SUSE Xen VM Provisioning Actions	23
1.10 Troubleshooting Linux KVM VM Provisioning Operations	25
2 Troubleshooting Cloud Manager Application Issues	27
2.1 Business Services Issues	27

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 Cloud Manager Orchestration Issues,” on page 9](#)
- ♦ [Chapter 2, “Troubleshooting Cloud Manager Application Issues,” on page 27](#)

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.

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 Troubleshooting Cloud Manager Orchestration Issues

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 9](#)
- ♦ [Section 1.2, “Troubleshooting Orchestration Server Issues,” on page 10](#)
- ♦ [Section 1.3, “Troubleshooting Orchestration Console Issues,” on page 12](#)
- ♦ [Section 1.4, “Troubleshooting Orchestration Agent Issues,” on page 13](#)
- ♦ [Section 1.5, “Troubleshooting General VM Management Issues,” on page 13](#)
- ♦ [Section 1.6, “Troubleshooting vSphere VM Provisioning Actions,” on page 18](#)
- ♦ [Section 1.7, “Troubleshooting Citrix Xen VM Provisioning Operations,” on page 21](#)
- ♦ [Section 1.8, “Troubleshooting Hyper-V VM Provisioning Operations,” on page 21](#)
- ♦ [Section 1.9, “Troubleshooting SUSE Xen VM Provisioning Actions,” on page 23](#)
- ♦ [Section 1.10, “Troubleshooting Linux KVM VM Provisioning Operations,” on page 25](#)

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 9](#)

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 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 10](#)
- ♦ [“Orchestration Server Might Hang if the System Clock Is Changed Abruptly” on page 10](#)
- ♦ [“Authentication to an Active Directory Server Might Fail” on page 10](#)
- ♦ [“The Orchestration Server Must Have Sufficient RAM” on page 11](#)
- ♦ [“Calling terminate\(\) from within a Job Class Allows the JDL Thread Execution to Continue” on page 11](#)
- ♦ [“Java programs That Use the JDL Exec Class Might Hang” on page 11](#)

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 2.1.5 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 2.1.5 Orchestration Installation Guide*.

Calling terminate() from within a Job Class Allows the JDL Thread Execution to Continue

Source: Cloud Manager Orchestration Server

Explanation: Calling `terminate()` from within the Job class does not immediately terminate the JDL thread of that job; instead, it sends a message to the server requesting termination of the job.

Action: This can take time to occur (because subjobs need to be recursively terminated and joblets cancelled), so if the calling JDL thread needs to terminate immediately, immediately follow the invocation of this method with `return`.

Java programs That Use the JDL Exec Class Might Hang

Source: Cloud Manager Orchestration Server

Explanation: Processes that are spawned by using the JDL Exec class on a Windows Orchestration Agent might hang when the spawned process attempts to read from `stdin`.

Action: To work around this issue, use the following steps to turn off the enhanced ExecWrapper:

1. In the Explorer tree of the Orchestration Console, select the job that you want to change.
2. In the admin view of the job, select the JDL Editor tab to open the JDL Editor.
3. Paste the following code into the editor:

```
e = Exec()  
e.setUseJvmRuntimeExec(True)
```

4. Save the changes.

NOTE: Disabling the enhanced ExecWrapper also makes other process control features provided as part of the ExecWrapper unavailable, such as running the process as a different user than the Orchestration Agent, or redirection of files (`Exec.setStdoutFile`, `Exec.setStderrFile` and `Exec.setStrinFile`).

For more information about the JDL Exec class, see the [Cloud Manager 2 JDL documentation \(http://www.novell.com/documentation/cloudmanager2/resources/jdljavadoc_2/com/novell/zos/jdl/Exec.html\)](http://www.novell.com/documentation/cloudmanager2/resources/jdljavadoc_2/com/novell/zos/jdl/Exec.html).

1.3 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 12](#)
- ♦ [“After Installing the Orchestration Agent on VM, the VM is Not Displayed as a Resource in the Orchestration Console” on page 13](#)

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.

1.4 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 13](#)

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.

1.5 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 14](#)

- ♦ “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 16
- ♦ “Admin Password on Windows 2003/2008 Workloads Cannot Be Set by Users” on page 16
- ♦ “Unable to Provision a VM to Another Cluster Node Due to Reason “VM Networks Are Not Available”” on page 17
- ♦ “If Multiple Workloads are Cloned Simultaneously, They are not Load Balanced Across Repositories” on page 17
- ♦ “Block Disks Show up as Regular Vdisks in Orchestrate” 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/rhel4tmp12/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>
```

```

value="never_use_this_resource"
this resource from consideration." />
factvalue="resource.provision.vmhost.cluster" />
factvalue="resource.provision.vmhost" />
fact="vmhost.resource.groups" value="another_vmhost_group"
not in the exclude_me group, when some_ boolean_fact is false,
some_other_boolean_fact is true, (but also not some other things),
in another_vmhost_group"/>

```

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.

Admin Password on Windows 2003/2008 Workloads Cannot Be Set by Users

Source: Windows 2003/2008 workloads in the Cloud Manager Web Console accessed by users.

Explanation: In order for a user to set the Administrator pass when configuring a Windows 2003/2008 workload, the VM template (from which the workload is created) must not have an Administrator password set.

Action: To leave the Administrator password unset on a VM template, you must turn off the complex password setting in the password policy.

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.

If Multiple Workloads are Cloned Simultaneously, They are not Load Balanced Across Repositories

Source: The Orchestration Server

Explanation: When multiple workloads are being cloned at the same time, the cloning process looks at the current state of the storage repositories to determine which repository should be used.

Possible Cause: With multiple asynchronous cloning processes running concurrently, the utilization of the repository does not reflect the state at the completion of other running clone processes. This leads to one repository being identified as the preferred repository until prior cloning jobs have finished running.

Action: Perform a single cloning operation at a time in order to achieve true load balancing, or be aware that multiple cloning operations can result in workload distributions between repositories that are not truly load balanced.

Block Disks Show up as Regular Vdisks in Orchestrate

Source: The Orchestration Client.

Explanation: The first time discovery is run after adding a new block device to a VM, the block device is marked as a regular vdisk in the repository.

Possible Cause: The VM discovery took place before the pdisk was discovered, and Orchestrator hasn't matched them up.

Action: Run the VM discovery process a second time.

1.6 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 18](#)
- ♦ [“\(503\) Service Unavailable Errors Might Occur while Cloning vSphere VMs” on page 18](#)
- ♦ [“Invalid Datastore Path Error” on page 19](#)
- ♦ [“Running Provisioning Operations on a Batch of vSphere VMs Results in JDL Event Handler Errors” on page 19](#)
- ♦ [“Moving a VM Host in vSphere Results in Duplicate Repositories” on page 20](#)

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()`.

Moving a VM Host in vSphere Results in Duplicate Repositories

Source: The Orchestration Server.

Explanation: If you move a VM host in your vSphere environment and then you subsequently perform a discovery in the Orchestration Console, the console displays duplicate repositories for the host that was moved.

Action: After you rediscover VM hosts and repositories in the Orchestration Console, you should delete the old repository grid object from the Explorer tree view in the Orchestration Console. Identify the repository to be deleted by checking the name of the datacenter, which is included in the `repository.datacenter` fact. If the value for this fact is the name of the old datacenter, this is the repository you want to delete.

1.7 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 21](#)

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.

1.8 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 21](#)
- ♦ [“Hyper-V Provisioning Jobs Fail When Several Jobs Are Started Simultaneously” on page 22](#)
- ♦ [“OS info for Hyper-V VMs is Not Always Auto-Discovered by the Orchestration Agent.” on page 22](#)
- ♦ [“Limitations of Linux VMs as Guests on Hyper-V” on page 22](#)

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 2.1.5 Orchestration Console Reference](#).

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.

Limitations of Linux VMs as Guests on Hyper-V

Source: Hyper-V provisioning adapter job

Explanation: The Orchestration Server does not support the *Create Template* or *Clone* actions for Linux-based Hyper-V VMs.

Action: None.

1.9 Troubleshooting SUSE 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 23](#)
- ♦ [“Multiple Instances of the Same Xen VM Running when Located on Shared Storage” on page 23](#)
- ♦ [“Running xm Commands on an Old Xen VM Host Causes Server to Hang” on page 24](#)

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
```

1.10 Troubleshooting Linux KVM VM Provisioning Operations

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

- ♦ [“The Save Config Action Does Not change the vnic.model Fact Value” on page 25](#)
- ♦ [“VM Migration Fails if Target Host is Unable to Resolve its Hostname” on page 25](#)

The Save Config Action Does Not change the vnic.model Fact Value

Source: Orchestration Server.

Explanation: The value of the `vnic.model` fact on a KVM VM is not set to “virtio” by default. This might cause a slowdown in the vNIC performance for that VM.

Changing the `vnic.model` fact value from “hypervisor default” to “virtio” in the Orchestration Console and then performing the *Save Config* action does not change the value.

Action: Set the model for the vNIC (`virtio`) in the KVM `virt-manager`, then rediscover the VM in the Orchestration Console.

VM Migration Fails if Target Host is Unable to Resolve its Hostname

Source: During a migration, the migration fails with a name resolution error.

Explanation: In order to complete the migration, the target host needs to be able to resolve its host name to an IP address.

Possible Cause: In some cases, a host’s name isn’t properly added to the `/etc/hosts` file.

Action: Update the `/etc/hosts` file to include an entry for the host’s name, or ensure that the host is properly able to resolve its own name via DNS.

2 Troubleshooting Cloud Manager Application Issues

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

2.1 Business Services Issues

The following sections provide solution to the problems you might encounter while using thCloud Manager:

- ♦ [“Business Service Workloads Remain in the Building or Provision state” on page 27](#)

Business Service Workloads Remain in the Building or Provision state

Source: Cloud Manager Application Server

Explanation: During the building phase of a new business service’s workload and the startup phase of a deployed workload, it is possible for the workload to be unable to be assigned to a host. This can occur when no hosts in the host group have the available resources to meet the workload resource requirements.

In the build phase, the business service remains in the *Building* state until a host becomes available for the workload. In the startup phase, the workload remains in the *Provision* state until a host becomes available.

Action: You have several options to resolve this issue:

- ♦ Shut down a workload to free up the required resources on a host. If possible, select a workload that can be restarted on another host.
- ♦ Add another host to the host group.

