

PlateSpin Orchestrate 2.0.2 Readme

Novell®

June 24, 2009

The information in this Readme file pertains to PlateSpin® Orchestrate from Novell®, the product that manages virtual resources and controls the entire life cycle of each virtual machine in the data center. PlateSpin Orchestrate also manages physical resources.

This document provides descriptions of limitations of the product or known issues and workarounds, when available. The issues included in this document were identified when PlateSpin Orchestrate 2.0 was initially released.

The information is current with the release of PlateSpin Orchestrate 2.0.2.

- ♦ [Section 1, “Readme Updates,” on page 1](#)
- ♦ [Section 2, “ZENworks Orchestrator Is Now PlateSpin Orchestrate,” on page 2](#)
- ♦ [Section 3, “Network File System Issues,” on page 2](#)
- ♦ [Section 4, “YaST Issues,” on page 3](#)
- ♦ [Section 5, “Installation Issues,” on page 3](#)
- ♦ [Section 6, “Upgrade Issues,” on page 4](#)
- ♦ [Section 7, “PlateSpin Orchestrate Server Issues,” on page 7](#)
- ♦ [Section 8, “PlateSpin Orchestrate Monitoring Issues,” on page 9](#)
- ♦ [Section 9, “PlateSpin Orchestrate Development Client Issues,” on page 10](#)
- ♦ [Section 10, “PlateSpin Orchestrate VM Client Issues,” on page 10](#)
- ♦ [Section 11, “Virtual Machine Management Issues,” on page 11](#)
- ♦ [Section 12, “Virtual Center VM Issues,” on page 14](#)
- ♦ [Section 13, “Xen VM Technology Issues,” on page 15](#)
- ♦ [Section 14, “NPIV Issues,” on page 16](#)
- ♦ [Section 15, “Documentation Conventions,” on page 17](#)
- ♦ [Section 16, “Legal Notices,” on page 17](#)

1 Readme Updates

The following information was added to or modified in the product readme at the release of PlateSpin Orchestrate 2.0.2. Obsolete information formerly included in the readme has been deleted.

- ♦ [Section 5.1, “Configuration Programs Do Not Include a Way to Edit the Agent Configuration,” on page 3](#)
- ♦ [Section 5.2, “The Orchestrate Configuration Program Has a Faulty Dependency When Installing the Agent on RHEL 5,” on page 4](#)

- ♦ Section 5.3, “Installing the Orchestrate Agent into a VM Considers Only the First Disk,” on page 4
- ♦ Section 5.4, “Uninstalling the Orchestrate Agent on Windows Machines Requires A Subsequent Reboot,” on page 4
- ♦ Section 6.6, “Audit Database Values Are Not Preserved In An Upgrade,” on page 7
- ♦ Section 7.3, “Orchestrate Server Might Hang if System Clock is Changed Abruptly,” on page 8
- ♦ Section 7.4, “Authentication to Active Directory Server Might Fail,” on page 8
- ♦ Section 7.5, “The vmhost.loadindex.slots Fact Is Not Dynamic,” on page 9
- ♦ Section 8.1, “Monitoring Agent on Windows 2008 Does Not Display Metrics on Web Interface,” on page 9
- ♦ Section 9.3, “Cloning a VM onto the datagrid repository (ZOS) is incorrectly available as an option in the Development Client,” on page 10
- ♦ Section 10.3, “Cannot increase the number of vCPUs on a running Xen VM,” on page 11
- ♦ Section 11.3, “VM State Is Not Properly Discovered for VMs Located on Shared Storage Repositories,” on page 12
- ♦ Section 11.4, “Canceling VM Build Fails on SLES 11 VM Host,” on page 13
- ♦ Section 11.5, “SUSE Linux VMs Might Attempt To Partition A Read-only Device,” on page 13
- ♦ Section 11.6, “The Setting for Gateway IP Address Might Be Confusing,” on page 14
- ♦ Section 11.7, “The Hyper-V Provisioning Adapter Sends Erroneous Status When Discovering VMs,” on page 14
- ♦ Section 12.4, “Virtual Center Discovery Jobs Run Only on Physical Resources,” on page 15
- ♦ Section 13.4, “Lock on a VM Protects only Against a Second VM from Provisioning,” on page 16

2 ZENworks Orchestrator Is Now PlateSpin Orchestrate

PlateSpin Orchestrate 2.0 is the latest product release of ZENworks® Orchestrator. The rebranded PlateSpin Orchestrate product is now part of the PlateSpin Workload Management product portfolio from Novell. For more details, see “[Product Rebranding](#)” in the *PlateSpin Orchestrate 2.0 Getting Started Reference*.

3 Network File System Issues

The following information is included in this section:

- ♦ Section 3.1, “Orchestrate Agent Fails to Set the UID on Files Copied from the Datagrid,” on page 2

3.1 Orchestrate Agent Fails to Set the UID on Files Copied from the Datagrid

If Network File System (NFS) is used to mount a shared volume across nodes that are running the Orchestrate 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.

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.

4 YaST Issues

The following information is included in this section:

- ♦ [Section 4.1, “YaST Uninstall Feature Is Not Supported,” on page 3](#)

4.1 YaST Uninstall Feature Is Not Supported

The uninstall feature in YaST and YaST2 is not supported in this release of PlateSpin Orchestrate.

5 Installation Issues

The following information about installation is included in this section:

- ♦ [Section 5.1, “Configuration Programs Do Not Include a Way to Edit the Agent Configuration,” on page 3](#)
- ♦ [Section 5.2, “The Orchestrate Configuration Program Has a Faulty Dependency When Installing the Agent on RHEL 5,” on page 4](#)
- ♦ [Section 5.3, “Installing the Orchestrate Agent into a VM Considers Only the First Disk,” on page 4](#)
- ♦ [Section 5.4, “Uninstalling the Orchestrate Agent on Windows Machines Requires A Subsequent Reboot,” on page 4](#)

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

Although the scenario is not supported in a production environment, it is common in demonstration or evaluation situations to install the PlateSpin Orchestrate Agent and the PlateSpin Orchestrate 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.

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

5.2 The Orchestrate Configuration Program Has a Faulty Dependency When Installing the Agent on RHEL 5

When you install the Orchestrate Agent on a RHEL 5 server, three packages are needed:

- ♦ RHEL5/novell-zenworks-zos-java-1.5.0_sun_update11-52.x86_64.rpm
- ♦ RHEL5/novell-zenworks-zos-agent-2.0.1-61007.i586.rpm
- ♦ RHEL5/novell-zenworks-orch-config-2.0.1-36.noarch.rpm

When you try to install RHEL5/novell-zenworks-orch-config-2.0.1-36.noarch.rpm, the installation fails. This is due to a faulty dependency on a python-xml package.

To work around this problem,

- 1 Verify you that libxml2-python package has been installed on your RHEL 5 system (install if not present).

```
rpm -q libxml2-python
```

- 2 Install the configuration package with the nodeps flag:

```
rpm -i --nodeps novell-zenworks-orch-config-2.0.2-36.noarch.rpm
```

5.3 Installing the Orchestrate Agent into a VM Considers Only the First Disk

If you choose to install the Orchestrate Agent into a VM, and if the VM has more than one associated disk, the vmprep job might fail, stating that it cannot recognize the installation target.

The vmprep job considers only the first listed VM disk when attempting to install the Orchestrate Agent. If the vmprep job cannot find the VM's operating system contained within the first disk, the job vmprep job fails. Therefore, when using a VM that has more than one associated disk, the disk image containing the OS install must always be listed first.

To work around the issue, reconfigure the VM so that the disk containing the OS install is listed first and any other VM disks are listed after it.

5.4 Uninstalling the Orchestrate Agent on Windows Machines Requires A Subsequent Reboot

If you install the Orchestrate Agent on a Windows machine (physical or virtual), then you uninstall the same agent and then install it again followed by a reboot, the agent does not start. This scenario occurs only if you manually installed the agent using the GUI installer originally.

To work around this problem, uninstall the agent, then reboot before installing the agent again using the installer.

6 Upgrade Issues

The following information is included in this section:

- ♦ [Section 6.1, “Upgrade Packages Are Labeled As Version 1.3,” on page 5](#)
- ♦ [Section 6.2, “Upgrading a Stopped Server Might Cause the Upgrade to Hang,” on page 6](#)

- ♦ Section 6.3, “Currently Defined Job Schedule Deployment States Are Overwritten on Upgrade,” on page 6
- ♦ Section 6.4, “OpenWBEM Users Need to Reconfigure after Upgrading,” on page 6
- ♦ Section 6.5, “Opening A Monitor Should Not Require Authentication,” on page 6
- ♦ Section 6.6, “Audit Database Values Are Not Preserved In An Upgrade,” on page 7

6.1 Upgrade Packages Are Labeled As Version 1.3

When you use either YaST or ZENworks Linux Management commands to upgrade ZENworks Orchestrator 1.3 components to PlateSpin Orchestrate 2.0 components on a server, the RPM packages listed after an upgrade () are labeled as version 1.3.

To work around this issue, you can use ZENworks Linux Management commands to upgrade the pattern versioning (depending on which patterns are installed) and install packages missed in the upgrade.

The following command adds packages not installed during the upgrade (avoiding an upgrade of older patterns, warehouse and orch_config) and adds the correct package version to the PlateSpin Orchestrate patterns already installed:

```
patterns_to_upgrade=$(rug --terse pt -i | grep zw | grep -v warehouse | grep -v orch_config | cut -d'|' -f2)
```

This command updates the pattern versioning:

```
rug in -t pattern $patterns_to_upgrade -y
```

IMPORTANT: For technical reasons, the orch-config package must remain at version 1.3. Do not remove this package.

Use the following command to remove the VM Warehouse package:

```
rug rm -t pattern zw_vm_warehouse -y
```

A system message is displayed:

```
The following packages will be removed:
zw_vm_warehouse 1.3-0 (system)
```

Run the following command to remove the remainder of the VM Warehouse component:

```
rug rm novell-zenworks-vmwarehouse-base novell-zenworks-vmwarehouse-
cimproviders -y
```

A system message is displayed:

```
The following packages will be removed:
novell-zenworks-vmwarehouse-base 1.3.0-46 (system)
novell-zenworks-vmwarehouse-cimproviders 1.3.0-27 (system)
```

6.2 Upgrading a Stopped Server Might Cause the Upgrade to Hang

If you use the standard command (`/etc/init.d/novell-zosserver stop`) to stop the PlateSpin Orchestrate prior to the upgrade, the preinstallation script detects that no snapshot was taken of the server, so it restarts the server and then stops it again to take a snapshot before upgrading the server package. If the grid has many objects, the `rug` command hangs during the upgrade process (that is, the `rug` command described in “[Upgrading ZENworks Orchestrator Server Packages at the Command Line](#)” in the *PlateSpin Orchestrate 2.0 Upgrade Guide*).

In order to execute a successful upgrade, we recommend that you keep the Orchestrate Server running during the upgrade or stop it by using the `--snapshot` flag (for example, `/etc/init.d/novell-zosserver stop --snapshot`) before the upgrade.

6.3 Currently Defined Job Schedule Deployment States Are Overwritten on Upgrade

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 ZENworks Orchestrator 1.3 to PlateSpin Orchestrate 2.0.

If you want to re-enable or disable a job after the upgrade, you need to open the Job Scheduler in the PlateSpin Orchestrate Development Client and manually change the deployment state.

For more information, see “[Creating or Modifying a Job Schedule](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*.

6.4 OpenWBEM Users Need to Reconfigure after Upgrading

Because PlateSpin Orchestrate 2.0 uses SFCB as the CIM provider for its VM Builder instead of using the OpenWBEM tool used in ZENworks Orchestrator 1.3, an upgrade from 1.3 to 2.0 causes a conflict for users who want to continue to run OpenWBEM. Both of these tools use the same port.

If you want to continue to use OpenWBEM, we recommend that you change the default TCP/IP port assigned to OpenWBEM and then make appropriate adjustments to any application that uses OpenWBEM.

6.5 Opening A Monitor Should Not Require Authentication

ZENworks Orchestrator 1.3 monitoring requires a username and password login. Although login is no longer required for opening a monitor (that is, `http://monitoring_server/monitor`) in PlateSpin Orchestrate 2.0.2, the Apache 2 Server continues to prompt for a login after the upgrade.

The reason for this login requirement after upgrade is that the `ganglia-auth.conf` file, created when the original installation configuration script ran, is not upgraded.

Because this version 1.3 file is necessary for the monitoring function, and because it is not updated with the other files in the rpm, you need to change it manually to be compatible with the 2.0.2 system.

The contents of the 1.3 are as follows:

```

Alias /monitor /var/opt/novell/zenworks/monitor/htdocs
<Directory /var/opt/novell/zenworks/monitor/htdocs>
    Order Allow,Deny
    Allow from all

    AuthType Basic
    AuthName Monitoring

    <IfDefine !DCA-NOLDAP>
        AuthBasicProvider ldap
    AuthLDAPURL ldaps:///
        AuthzLDAPAuthoritative on

require ldap-group
</IfDefine>

<IfDefine DCA-NOLDAP>
    AuthBasicProvider file
    AuthUserFile /opt/novell/zenworks/monitor/web/user.dat
    AuthGroupFile /opt/novell/zenworks/monitor/web/grpfile.dat

    require group DCA-monitoring
</IfDefine>
</Directory>

```

You need to replace this 1.3 information with the following 2.0.2 information:

```

Alias /monitor /var/opt/novell/zenworks/monitor/htdocs
<Directory /var/opt/novell/zenworks/monitor/htdocs>
    Order Allow,Deny
    Allow from all
</Directory>

```

This content is compatible with the 2.0.2 Orchestrate system.

When you have made the change, you need to restart the Apache 2 server.

6.6 Audit Database Values Are Not Preserved In An Upgrade

If you upgrade the PlateSpin Orchestrate Server to version 2.0.x, the following values for the audit database configuration are not preserved in order to maintain security:

- ♦ JDBC connection URL (including the previously-defined database name)
- ♦ previously-specified database username
- ♦ previously-specified database password

The administrator is responsible to know the audit database owner username and password and to enter them during the upgrade process.

7 PlateSpin Orchestrate Server Issues

The following information is included in this section:

- ♦ [Section 7.1, “Administration Clients Can’t Log In to an Orchestrate Server without DNS,” on page 8](#)

- [Section 7.2, “Orchestrate Server Might Appear to Be Deadlocked when Provisioning Large Numbers of Jobs with Subjobs,” on page 8](#)
- [Section 7.3, “Orchestrate Server Might Hang if System Clock is Changed Abruptly,” on page 8](#)
- [Section 7.4, “Authentication to Active Directory Server Might Fail,” on page 8](#)
- [Section 7.5, “The vmhost.loadindex.slots Fact Is Not Dynamic,” on page 9](#)

7.1 Administration Clients Can’t Log In to an Orchestrate Server without DNS

The administration clients, including the CLI tool, the Orchestrate Development Client, and the Orchestrate VM Client, communicate with the Orchestrate Server over RMI. RMI requires a correctly configured DNS service.

If you encounter connectivity problems (that is, you are unable to remotely log in to the server through these clients), verify your DNS configuration for proper hostname resolution.

7.2 Orchestrate Server Might Appear to Be Deadlocked when Provisioning Large Numbers of Jobs with Subjobs

In some deployments where a large number of running jobs spawn subjobs, the running jobs might appear to stop, leaving jobs in the queue. This occurs because of job limits set in the Orchestrate Server to avoid overload or “runaway” conditions.

If this deadlock occurs, you can slowly adjust the job limits to tune them according to your deployment. For more information, see “[job.limits](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*.

7.3 Orchestrate Server Might Hang if System Clock is Changed Abruptly

As with many applications, you should avoid abrupt changes in the system clock on the machine where the PlateSpin Orchestrate 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.

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.

7.4 Authentication to Active Directory Server Might Fail

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).

The checking logic in the current AuthLDAP auth provider assumes that if any `namingContext` entry is returned, then it has found the domain and it stops searching. If you encounter this issue, you need to manually configure LDAP as a generic LDAP server, which offers many more configuration options.

7.5 The vmhost.loadindex.slots Fact Is Not Dynamic

If a VM is waiting to be provisioned to a VM host and a host is not available, you might consider editing the `vmhost.maxvmslots` fact to fix it (`vmhost.loadindex.slots` is the calculation of `vmhost.maxvmslots` divided by `vmhost.vm.instanceids`), but doing so has no effect.

Changing the `maxvmslots` fact is not considered in the logic until an unrelated provision occurs on that VM host (or until the current provision can occur because, for instance, an original slot becomes available).

Future releases of the product will have updated host selection logic to re-read the `maxvmslots` fact when it dynamically changes.

8 PlateSpin Orchestrate Monitoring Issues

The following information is included in this section:

- ♦ [Section 8.1, “Monitoring Agent on Windows 2008 Does Not Display Metrics on Web Interface,” on page 9](#)

8.1 Monitoring Agent on Windows 2008 Does Not Display Metrics on Web Interface

Beginning with the PlateSpin Orchestrate 2.0.2 release, the ganglia monitoring agent (`gmond`) installed on a Windows server (including Windows 2008) supplies metrics information to the Monitoring Server. To make this data appear in the Monitoring Server’s Web service (`http://ip_or_dns_of_monitoring_server/monitor`) you need to use Wordpad to edit the `gmond.conf` file, adding the following two entries:

```
udp_send_channel {
  host = 127.0.0.1
  port = 8649
}

udp_send_channel {
  host = 10.201.10.1
  port = 8649
}
```

The IP address in the second entry (shown as `10.201.10.1` in the example above) is the IP address of the Orchestrate Monitoring Server (where `gmetad` is installed). You can also use the DNS name.

You need to keep the `127.0.0.1` to send metric information to the Orchestrate Agent so that the `resource.metric` facts become populated for that resource.

You can also change the following entry in `gmond.conf` from `localhost` to something more recognizable or intuitive. For example

```
host {
  location = "localhost"
}
```

changed to

```
host {
  location = "pe2950-10-1"
}
```

To troubleshoot the edited file, run the `gmond.exe` file with the appropriate parameters and options, as in the following example:

```
C:\Program Files\Gmond>Gmond.exe -c c:\gmond.conf -d5
```

9 PlateSpin Orchestrate Development Client Issues

The following information is included in this section:

- ♦ [Section 9.1, “Using the Orchestrate Development Client in a Firewall Environment,” on page 10](#)
- ♦ [Section 9.2, “Windows Sysprep Config Is Nonfunctional,” on page 10](#)
- ♦ [Section 9.3, “Cloning a VM onto the datagrid repository \(ZOS\) is incorrectly available as an option in the Development Client,” on page 10](#)

9.1 Using the Orchestrate Development Client in a Firewall Environment

Using the PlateSpin Orchestrate Development Client in a firewall environment (NAT, in particular) is not supported for this release. The Orchestrate Development Client uses RMI to communicate with the server, and RMI connects to the initiator on dynamically chosen port numbers. To use the Development Client in a firewall environment, you need to use a remote desktop or VPN product.

9.2 Windows Sysprep Config Is Nonfunctional

Selecting a VM Object in the Explorer tree of the PlateSpin Orchestrate Development Client opens a properties page in the workspace area of the client. One panel of this page, titled *Windows Sysprep Config* has configuration settings that are currently nonfunctional. The feature is not implemented.

9.3 Cloning a VM onto the datagrid repository (ZOS) is incorrectly available as an option in the Development Client

The datagrid (ZOS) repository is not supported as a cloning target. However, it is listed in the Development Client as an option to select when cloning a new VM from a template.

In the VM Client, the ZOS repository is not presented as an option when cloning.

To work around this issue, do not select the ZOS option when cloning.

10 PlateSpin Orchestrate VM Client Issues

The following information is included in this section:

- ♦ [Section 10.1, “VM Does Not Start When a Local Repository Is Assigned to Multiple Hosts,” on page 11](#)
- ♦ [Section 10.2, “Not Configuring a Display Driver Triggers a Pop-Up Message,” on page 11](#)
- ♦ [Section 10.3, “Cannot increase the number of vCPUs on a running Xen VM,” on page 11](#)

10.1 VM Does Not Start When a Local Repository Is Assigned to Multiple Hosts

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 any 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 will not start and no longer display in the VM Client or Development Client, making it impossible to recover.

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.

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.

10.2 Not Configuring a Display Driver Triggers a Pop-Up Message

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

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.

10.3 Cannot increase the number of vCPUs on a running Xen VM

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.

Workaround: 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.

11 Virtual Machine Management Issues

The following information is included in this section:

- ♦ [Section 11.1, “Autoprep When Using LVM as a Volume Manager,” on page 12](#)
- ♦ [Section 11.2, “Manually Created VM Might Display “Under Construction” on the VM Icon,” on page 12](#)

- ♦ [Section 11.3, “VM State Is Not Properly Discovered for VMs Located on Shared Storage Repositories,” on page 12](#)
- ♦ [Section 11.4, “Canceling VM Build Fails on SLES 11 VM Host,” on page 13](#)
- ♦ [Section 11.5, “SUSE Linux VMs Might Attempt To Partition A Read-only Device,” on page 13](#)
- ♦ [Section 11.6, “The Setting for Gateway IP Address Might Be Confusing,” on page 14](#)
- ♦ [Section 11.7, “The Hyper-V Provisioning Adapter Sends Erroneous Status When Discovering VMs,” on page 14](#)

11.1 Autoprep When Using LVM as a Volume Manager

If you plan to prepare Virtual Machines that use LVM as their volume manager on a SLES VM Host, and if that VM Host also uses LVM as its volume manager, you cannot perform the autoprep if the VM has an LVM volume with the same name as one already mounted on the VM Host. This is because LVM on the VM Host can mount only one volume with the same name.

To work around this issue, ensure that the volume names on the VM Hosts and Virtual Machines are different.

11.2 Manually Created VM Might Display “Under Construction” on the VM Icon

If you manually install the Orchestrate Agent on a running VM for which there is a corresponding VM Grid Object, you must use the same name for the agent and for the Grid Object of the VM that contains the agent. If different names are used, an Under Construction flag overlays the VM icon in the Orchestrate Development Client.

This flag is used in constraints to prevent the attempted provisioning of a VM that is not yet built or that is not completely set up. The flag is cleared automatically by the provisioning adapters when names match.

If the names do not match, you need to clear the flag by manually adjusting the `agent.properties` file to match the names or by reinstalling the Orchestrate Agent on the VM and making sure the names match.

11.3 VM State Is Not Properly Discovered for VMs Located on Shared Storage Repositories

When a VM is provisioned to run, the name of the VM host selected is placed in the `resource.provisioner.recommendedhost` fact on the VM grid object. Whenever a “resync state” action is issued on the VM, the provisioning adapter queries only the VM host that is listed in the `resource.provisioner.recommendedhost` fact. In some instances, it is possible that the VM has actually moved and is no longer running on that host. In this case, the provisioning adapter incorrectly reports a status of “down,” even though the VM is actually running on another host.

For example, in a scenario where there are two VM hosts, “foo” and “bar,” both connected to shared storage, the sequence of events happens like this:

1. The administrator uses PlateSpin Orchestrate to provision a VM on the shared storage to the VM host “foo.” The `resource.provisioner.recommendedhost` fact is set equal to “foo.”

2. The VM on VM host “foo” is moved or stopped and then restarted without the knowledge of PlateSpin Orchestrate. This VM might now be running on a different VM host, for example, “bar.”
3. The “resynch state” action on the VM queries only the “foo” VM host, because that is the location where the VM was last known as running.
4. The VM state is reported as “down,” even though the VM is actually running on a different VM host: “bar.”

To work around this issue you can use one of two methods:

- ♦ Manually repeat the VM discovery on each VM host. The discovery job checks for any running VMs. When the VM is found to be running on another host, the VM state is corrected and the `resource.provisioner.recommendedhost` fact is reset to the actual host where the VM is running.
- ♦ Restart the Orchestrate Server (`/etc/init.d/novell-zosserver restart`). This causes all the Orchestrate Agents to log in to the Orchestrate Server again. The restart also triggers the same VM host discovery job on all VM hosts as described in first workaround method above.

11.4 Canceling VM Build Fails on SLES 11 VM Host

If you attempt to cancel a VM build already in progress on a SLES 11 VM host, the VM build job might fail to cancel the running VM build, leaving the VM running on the VM host. The behavior occurs when canceling either from the Orchestrate Development Client or the Orchestrate VM Client.

To work around the issue, cancel the build job normally from either client, log into the physical machine where the VM has been building, and then manually destroy the VM (for example, by using the `xm destroy` command). Afterward, you need to manually resync the VM grid object state by using either the Orchestrate Development Client or the Orchestrate VM Client.

11.5 SUSE Linux VMs Might Attempt To Partition A Read-only Device

When building a SUSE Linux VM and specifying 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.

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 most available disk space without considering the other device attributes. For example, if a specified CD-ROM device happens to be larger in size than the specified hard disk device, YaST attempts to partition the CD-ROM device, which causes the VM installation to fail.

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.

11.6 The Setting for Gateway IP Address Might Be Confusing

Currently, the *Gateway IP Address* setting under the *Autoprep Network Adapter 0* and the Autoprep Network Adapter 1 sections of the *Info/Groups* tab for a VM object and VM template object is available in a list box.

Because the VM OS accepts only one default gateway, it accepts only the first setting in the list as the actual Gateway IP Address. The other settings are ignored.

11.7 The Hyper-V Provisioning Adapter Sends Erroneous Status When Discovering VMs

When the Hyper-V provisioning adapter runs a discovery for VM repositories (following the installation of the Orchestrate Agent), Java exceptions from the VM Manager appear in the Orchestrate Server log:

```
java.lang.ClassCastException
```

The Hyper-V provision adapter sends this message to the server erroneously. The message is ignored by the VM Manager.

12 Virtual Center VM Issues

The following information is included in this section:

- ♦ [Section 12.1, “Customization of Network Configuration Is Not Working Properly for VCenter 2.x Managed Virtual Machines,” on page 14](#)
- ♦ [Section 12.2, “Moving VM Templates across VCenter 2.x Managed Virtual Machines Hosts Is Not Supported,” on page 14](#)
- ♦ [Section 12.3, “Orchestrate Groups Do Not Match VMware Virtual Center Groups,” on page 15](#)
- ♦ [Section 12.4, “Virtual Center Discovery Jobs Run Only on Physical Resources,” on page 15](#)

12.1 Customization of Network Configuration Is Not Working Properly for VCenter 2.x Managed Virtual Machines

If you try to change the network configurations of virtual machine through the VCenter provisioning adapter, the settings might not be applied on virtual machines managed by VCenter 2.5 and VCenter 2.0.1

12.2 Moving VM Templates across VCenter 2.x Managed Virtual Machines Hosts Is Not Supported

The VM templates cannot be moved across VM hosts by using the *Move VM Template* option through the VCenter provisioning adapter. VCenter 2.x does not allow moving VM templates across VM hosts.

12.3 Orchestrate Groups Do Not Match VMware Virtual Center Groups

VMware* Virtual Center 2.0 introduced a new object grouping for clustering that is not supported by the VMware API adapter currently shipping with PlateSpin Orchestrate 2.0. The additional group might cause VM and host object mismatches after the PlateSpin Orchestrate system discovers VM images in VMware Virtual Center 2.0 and when you try to provision a VM in the cluster grouping.

To work around the issue, manually create a resource group in PlateSpin Orchestrate to match the group existing in Virtual Center 2.0. After you create the group to match the *cluster* group, you need to add the discovered resource to the new group and also add the new group to the *Available VM Resource Groups* for the VM Hosts that are to be used for provisioning the resource.

12.4 Virtual Center Discovery Jobs Run Only on Physical Resources

Currently, the vcenter2x discovery jobs in PlateSpin Orchestrate are constrained to run only on physical resources, not on VMs.

To work around this problem, you need to ensure that the Orchestrate Server views the resource where Virtual Center Server is running as a physical resource. Do not select the *Installing on VM* option when installing the agent.

13 Xen VM Technology Issues

The following information is included in this section:

- ♦ [Section 13.1, “Some Features Are Not Supported in the Xen Hypervisor,” on page 15](#)
- ♦ [Section 13.2, “Suspending A Fully Virtualized VM Makes VM Unrecoverable,” on page 15](#)
- ♦ [Section 13.3, “Running xm Commands on An Old Xen VM Host Causes Server to Hang,” on page 16](#)
- ♦ [Section 13.4, “Lock on a VM Protects only Against a Second VM from Provisioning,” on page 16](#)

13.1 Some Features Are Not Supported in the Xen Hypervisor

The “checkpoint” and “restore” features are not available for the Xen* provisioning adapter.

13.2 Suspending A Fully Virtualized VM Makes VM Unrecoverable

When suspending a 32-bit fully virtualized SLES 10 SP2 VM on a 64-bit host, Xen might put the VM into an unrecoverable state that prevents freeing the loopback device, starting the Virtual Machine, or deleting the VM from the Xen host. The loopback device can be freed up only by restarting the physical machine.

This is a known Xen problem when the paravirtualized drivers are installed on the fully virtualized machine.

To work around this problem, remove the paravirtualized driver from the fully virtualized machine by logging into the fully virtualized machine and removing the following package:

```
xen-kmp-default-3.2.0_16718_14_2.6.16.60_0.21-0.4
```

13.3 Running xm Commands on An Old Xen VM Host Causes Server to Hang

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

Make sure the Xen VM host has the latest Xen tools available by running the following command:

```
rpm -qa | grep xen-tools
```

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

```
Xen 3.3.1_18546_14
```

13.4 Lock on a VM Protects only Against a Second VM from Provisioning

When VM locking has been enabled and a Xen VM is running on a node and that node loses network connectivity to the Orchestrate Server, then a reprovisioning of the VM fails due to the lock, protecting the VM's image. The VM Client indicates that the VM is down, even though the VM might still be running on the node that has been cut off.

The failed reprovisioning sends a message to the VM Client informing the user about this situation:

```
The VM is locked and appears to be running on <host>
```

The error is added to the provisioner log.

Currently, the locks protect only against a second provisioning of the VM, not against moving the VM's image to another location. It is therefore possible to move the VM (because PlateSpin Orchestrate detects that the VM is down) and to reprovision it on another VM host.

If the original VM is still running on the cut-off VM host, this provisioning operation makes the VM crash. We recommend that you do not move the image, because unpurged, OS-level cached image settings might still exist.

14 NPIV Issues

The following issues and limitations have been identified with the NPIV disks:

- ♦ [Section 14.1, “Do Not Select the Moveable Option for NPIV disks,” on page 16](#)
- ♦ [Section 14.2, “Size of the NPIV Disk Is Always Zero,” on page 17](#)

14.1 Do Not Select the Moveable Option for NPIV disks

Do not select the *moveable* option for the NPIV disks in the Virtual Disk Editor. By default, it is not selected when you add a new NPIV disk.

14.2 Size of the NPIV Disk Is Always Zero

Do not configure the size option for the NPIV disk in the Virtual Disk Editor. The option is not applicable for the NPIV disks and the value is always displayed as zero (0).

15 Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions required when navigating menus in a user interface.

A trademark symbol (® , ™, etc.) denotes a Novell trademark; an asterisk (*) denotes a third-party trademark.

16 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. Please refer to <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 © 2008-2009 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. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.novell.com/company/legal/patents> (<http://www.novell.com/company/legal/patents>) and one or more additional patents or pending patent applications in the U.S. and in other countries.

For a list of 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>) at <http://www.novell.com/company/legal/trademarks/tmlist.html>.

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

