

# PlateSpin Orchestrate 2.5 Readme

March 29, 2010

Novell®

The information in this Readme file pertains to PlateSpin Orchestrate, 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 2.5 was initially released.

- ◆ [Section 1, “Readme Updates,” on page 1](#)
- ◆ [Section 2, “Network File System Issues,” on page 2](#)
- ◆ [Section 3, “YaST Issues,” on page 2](#)
- ◆ [Section 4, “Installation Issues,” on page 3](#)
- ◆ [Section 5, “Upgrade Issues,” on page 3](#)
- ◆ [Section 6, “PlateSpin Orchestrate Server Issues,” on page 5](#)
- ◆ [Section 7, “PlateSpin Orchestrate Monitoring Issues,” on page 10](#)
- ◆ [Section 8, “PlateSpin Orchestrate Development Client Issues,” on page 11](#)
- ◆ [Section 9, “PlateSpin Orchestrate VM Client Issues,” on page 13](#)
- ◆ [Section 10, “Virtual Machine Management General Issues,” on page 15](#)
- ◆ [Section 11, “vSphere VM Issues in the Development Client,” on page 19](#)
- ◆ [Section 12, “Xen VM Issues in the Development Client,” on page 22](#)
- ◆ [Section 13, “Hyper-V VM Issues in the Development Client,” on page 27](#)
- ◆ [Section 14, “Legal Notices,” on page 28](#)

## 1 Readme Updates

As PlateSpin Orchestrate 2.5 is deployed by Novell customers, issues of interest to all users are occasionally discovered and reported. This section includes such content, included after the initial release of the readme. Content is dated to direct your attention to newer items.

- ◆ [Section 1.1, “March 29, 2011,” on page 1](#)
- ◆ [Section 1.2, “October 15, 2010,” on page 2](#)

### 1.1 March 29, 2011

The following content items were added with this update:

- ◆ [Section 10.8, “Policies Applied to VM Resources Are Deleted,” on page 17](#)
- ◆ [Section 10.9, “VMs Provisioned from a VM Template Are Not Restarted When a VM Host Crashes,” on page 19](#)

## 1.2 October 15, 2010

The following content items were added with this update:

- ♦ [Section 5.6, “With \\$ Character in Password, the Administrator is Locked Out After Upgrade,” on page 5](#)
- ♦ [Section 12.15, “A Xen VM Cannot Provision If It Uses a Physical Disk Without a Repository,” on page 26](#)

## 2 Network File System Issues

The following information is included in this section:

- ♦ [Section 2.1, “Orchestrate Agent Fails to Set the UID on Files Copied from the Datagrid,” on page 2](#)

### 2.1 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.

---

## 3 YaST Issues

The following information is included in this section:

- ♦ [Section 3.1, “YaST Uninstall Feature Is Not Supported,” on page 2](#)

### 3.1 YaST Uninstall Feature Is Not Supported

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

## 4 Installation Issues

The following information about installation is included in this section:

- ♦ [Section 4.1, “Configuration Programs Do Not Include a Way to Edit the Agent Configuration,” on page 3](#)
- ♦ [Section 4.2, “Configuring the Orchestrate Agent on RHEL 4 Machines Fails,” on page 3](#)

### 4.1 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.

### 4.2 Configuring the Orchestrate Agent on RHEL 4 Machines Fails

If you install the Orchestrate Agent on a RHEL 4 machine and then try to configure it with PlateSpin Orchestrate 2.5 RPMs, the configuration script fails. This occurs because of a dependency on a Python 2.4 subprocess module, which is not included with RHEL 4.

To work around the problem, do one of the following:

- ♦ Remove the configuration RPMs for RHEL 4 and configure the agent manually by editing the `/opt/novell/zenworks/zos/agent/agent.properties` file.
- ♦ If the resource where you want to install the agent is a VM, use the *Install Agent* action available in the Development Client.
- ♦ Download and install the RPM that provides Python 2.4 support in RHEL 4. This file is available for download at the [Python download site \(http://www.python.org/download/releases/2.4/rpms/\)](http://www.python.org/download/releases/2.4/rpms/).

## 5 Upgrade Issues

The following information is included in this section:

- ♦ [Section 5.1, “Upgrading a Stopped Server Might Cause the Upgrade to Hang,” on page 4](#)
- ♦ [Section 5.2, “Currently Defined Job Schedule Deployment States Are Overwritten on Upgrade,” on page 4](#)
- ♦ [Section 5.3, “Audit Database Values Are Not Preserved in an Upgrade,” on page 4](#)
- ♦ [Section 5.4, “Upgrading vSphere from Orchestrate 2.0.2 to Orchestrate 2.5,” on page 4](#)

- ♦ [Section 5.5, “A Clone Does Not Inherit the Policy Associations of Its Upgraded Parent VM Template,”](#) on page 5
- ♦ [Section 5.6, “With \\$ Character in Password, the Administrator is Locked Out After Upgrade,”](#) on page 5

## 5.1 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 PlateSpin Orchestrate Server Packages at the Command Line”](#) in the *PlateSpin Orchestrate 2.5 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.

## 5.2 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 PlateSpin Orchestrate 2.0 or 2.1 to PlateSpin Orchestrate 2.5.

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 2.5 Development Client Reference*.

## 5.3 Audit Database Values Are Not Preserved in an Upgrade

If you upgrade the PlateSpin Orchestrate Server to version 2.5, 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.

## 5.4 Upgrading vSphere from Orchestrate 2.0.2 to Orchestrate 2.5

Although PlateSpin Orchestrate 2.0.2 required the [modification of `vsphere.policy`](http://www.novell.com/documentation/beta/ps_orchestrate21/ps_21_vmmgmt/data/bomts4t.html) ([http://www.novell.com/documentation/beta/ps\\_orchestrate21/ps\\_21\\_vmmgmt/data/bomts4t.html](http://www.novell.com/documentation/beta/ps_orchestrate21/ps_21_vmmgmt/data/bomts4t.html)) to include authentication credentials, this is no longer the case with PlateSpin Orchestrate 2.5.

In PlateSpin Orchestrate 2.5, the `vsphere.policy` file contains only the name of the credential from the Orchestrate Credential Manager for use in accessing vSphere. Instead of adding the credentials to the policy, you now need to create the credential using the Credential Manager. Then in `vsphere.policy`, you uncomment the `webservice_credential_name` item in the `vcenters` fact and set the value to the name of the newly created credential.

For more information, see [Step 5](#) in the “[Configuring the VMware vSphere Provisioning Adapter](#)” section of the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

## 5.5 A Clone Does Not Inherit the Policy Associations of Its Upgraded Parent VM Template

When the PlateSpin Orchestrate Server is upgraded, the parent-template/clone relationship is not re-created properly: clones do not inherit the policy associations that were created on the parent template.

Currently, it is not possible to modify policy associations on a cloned VM in PlateSpin Orchestrate, so if the cloned VM requires these associations, you can delete it in the Development Client, then rediscover it. After the discovery, you can apply the policies you want to this VM.

## 5.6 With \$ Character in Password, the Administrator is Locked Out After Upgrade

If the PlateSpin Orchestrate administrator used a dollar sign (\$) in his or her password in Orchestrate version 2.0.x, that character causes a lockout for administrator login after Orchestrate is upgraded to version 2.5.

To work around the issue, you should change the password to exclude the dollar sign.

The issue will be fixed in version 2.6 of PlateSpin Orchestrate.

# 6 PlateSpin Orchestrate Server Issues

The following information is included in this section:

- ◆ [Section 6.1, “Orchestrate Server Might Appear to Be Deadlocked When Provisioning Large Numbers of Jobs with Subjobs,” on page 6](#)
- ◆ [Section 6.2, “Orchestrate Server Might Hang if the System Clock Is Changed Abruptly,” on page 6](#)
- ◆ [Section 6.3, “Authentication to an Active Directory Server Might Fail,” on page 6](#)
- ◆ [Section 6.4, “Orchestrate Server Does Not Write to the Audit Database When the Maximum Queue Size Is Reached,” on page 6](#)
- ◆ [Section 6.5, “Orchestrate Server Might Shut Down When Managing Large Numbers of VMs and Resources,” on page 7](#)
- ◆ [Section 6.6, “Custom JDL Requires an “Import Sys” Directive,” on page 8](#)
- ◆ [Section 6.7, “The Orchestrate Server Might Require Caching of Computed Facts in a Grid with Large Numbers of Resources,” on page 8](#)
- ◆ [Section 6.8, “The Orchestrate Server Must Have Sufficient RAM,” on page 9](#)

- ♦ [Section 6.9, “Calling terminate\(\) from within a Job Class Allows the JDL Thread Execution to Continue,”](#) on page 9
- ♦ [Section 6.10, “Deploying Components Might Fail Intermittently,”](#) on page 9
- ♦ [Section 6.11, “Some Python Attributes Cannot Be Set on Job Objects in the JDL,”](#) on page 9

## 6.1 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 “[The Job Limits Panel](#)” in the *PlateSpin Orchestrate 2.5 Development Client Reference*.

## 6.2 Orchestrate Server Might Hang if the 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.

## 6.3 Authentication to an 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, 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.

## 6.4 Orchestrate Server Does Not Write to the Audit Database When the Maximum Queue Size Is Reached

A large number of audit database transactions occurring in a large grid might cause the audit database to become overloaded and lose the audit information. You can recognize this behavior by messages in the server log like the one below, or if you notice that data are missing in the audit database.

```
01.21 18:26:26: Audit,NOTICE: Object 10592 Pause will not not be written to
the
database because the queue size has reached its max: 200
01.21 18:26:26: Audit,NOTICE: Object 10593 Pause will not not be written to
the
database because the queue size has reached its max: 200
01.21 18:26:26: Audit,NOTICE: Object c-S-M2-14-01-21-13-35-41-200 will not not
be written to the database because the queue size has reached its max: 200
```

If you notice that the queue size has reached its maximum, increase the database buffer size of the audit database. The default size is current set at 200. At a minimum, you should set the value larger than the number of managed nodes.

To increase the queue size to 1000 records, for example, run the following command after authenticating with the `zosadmin login` command.

```
zosadmin set --mbean="local:facility=audit" --attr=QueueSizeMax --
type=Integer --value=1000
```

After you run this command, you must immediately restart the server:

```
rcnovell-zosserver restart
```

## 6.5 Orchestrate Server Might Shut Down When Managing Large Numbers of VMs and Resources

The PlateSpin Orchestrate system has been tested to a support level of 1000 VMs and 124 separate VM hosts being managed.

If these support levels are exceeded, the Orchestrate service (`novell-zosserver`) might shut down with the following log entry recorded in `/var/opt/novell/zenworks/zos/server/logs/server.log`:

```
ERROR: Out of Memory
ERROR : You might want to try the -mx flag to increase heap size.
```

To change the heap size:

- 1 From a text editor, open `/etc/init.d/novell-zosserver`.

- 2 Edit the start parameters in the file to increase heap size:

- 2a Change the following entry:

```
$ZOS_BIN start -d $ZOS_CONFIG > /dev/null
```

to

```
$ZOS_BIN start --jvmargs=-Xmx4g -d $ZOS_CONFIG > /dev/null
```

- 2b Save the file, then restart the server.

See “[Changing Orchestrate Server Default Parameters and Values](#)” in the *PlateSpin Orchestrate 2.5 Administrator Reference* for a list of attributes that can be adjusted to increase server performance under heavy load.

## 6.6 Custom JDL Requires an “Import Sys” Directive

If you utilize the `sys` module in custom JDL files that you use in PlateSpin Orchestrate, an “import `sys`” directive must be included in the file in the appropriate place. In prior versions of PlateSpin Orchestrate, you were not required to explicitly import the `sys` module, but this has changed in version 2.5.

In PlateSpin Orchestrate 2.5, if this import is not performed, you see the following error message:

```
NameError: name 'sys' is not defined
```

## 6.7 The Orchestrate Server Might Require Caching of Computed Facts in a Grid with Large Numbers of Resources

If your Orchestrate grid includes a large number of resources with associated Computed Facts, it is likely that these computed facts are evaluated with each Ready for Work message received by the broker from the Orchestrate Agent. These evaluations can cause an excessive load on the Orchestrate Server, causing a decrease in performance. You might see warnings in the server log similar to the following:

```
07.07 18:27:54: Broker,WARNING: ----- Long scheduler cycle time detected -----
07.07 18:27:54: Broker,WARNING: Total:3204ms, JDL thrds:8, TooLong:false
07.07 18:27:54: Broker,WARNING: Allocate:0ms [P1:0,P2:0,P3:0], Big:488
07.07 18:27:54: Broker,WARNING: Provision:4ms [P1:0,P2:0,P3:0], Big:253
07.07 18:27:54: Broker,WARNING: Msgs:3204ms [50 msg, max: 3056ms (3:RFW)]
07.07 18:27:54: Broker,WARNING: Workflow:[Timeout:0ms, Stop:0ms]
07.07 18:27:54: Broker,WARNING: Line:0ms, Preemption:0ms, (Big: 3), Mem:0ms
07.07 18:27:54: Broker,WARNING: Jobs:15/0/16, Contracts:10, AvailNodes:628
07.07 18:27:54: Broker,WARNING: PermGen: Usage [214Mb] Max [2048Mb] Peak
[543Mb]
07.07 18:27:54: Broker,WARNING: Memory: free [1555Mb] max [3640Mb]
07.07 18:27:54: Broker,WARNING: Msgs:483/50000 (recv:128692,sent:14202),
More:true
07.07 18:27:54: Broker,WARNING: -----
```


To work around this issue, we recommend that you cache the Computed Facts.

- 1 In the Explorer tree of the Orchestrate Development Client, expand the Computed Facts object, then select `vmbuilderPXEBoot`.

The `vmbuilderPXEBoot` fact does not change, so setting the cache here is safe from any automatic modifications.

- 2 In the Computed Facts admin view, select the *Attributes* tab to open the Attributes page.
- 3 In the Attributes page, select the *Cache Result for* check box, then in the newly active field, enter 10 minutes (remember to change the drop-down list to indicate *Minutes*).

This value must be greater than the default of 30 seconds.

- 4 Click the Save icon  to save the new configuration.

---

**NOTE:** If necessary, you can also cache other computed facts to improve server performance.

---



## 6.8 The Orchestrate Server Must Have Sufficient RAM

If the PlateSpin Orchestrate 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. The PlateSpin Orchestrate Server requires 3 GB of RAM to function with the preset defaults. 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 [Step 10](#) of the “[Installation and Configuration Steps](#)” in the *PlateSpin Orchestrate 2.5 Installation and Configuration Guide*. You can find similar information in [Section 6.5, “Orchestrate Server Might Shut Down When Managing Large Numbers of VMs and Resources,”](#) on page 7.

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

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. 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`.

## 6.10 Deploying Components Might Fail Intermittently

It is possible that when you attempt to deploy a component such as `.job`, `sjob`, `jdl`, `cfact`, `event`, `metric`, `policy`, `eaf`, `sar`, `sched`, `trig`, `python`, `pylib`; where prepackaged components are located in the `/opt/novell/zenworks/zos/server/components` directory, PlateSpin Orchestrate might intermittently fail the deployment, displaying a message similar to the following:

```
ERROR: Failed to deploy ./mem_free.<component> to <name_of_server>
      : TAP manager could not register
zos:deployer/<component>:com.novell.zos.facility.DefaultTAPManagerException:
Cannot locate zos:deployer/<component> in load status.
```

To work around this issue, restart the Orchestrate Server to bring the deployer back online.

## 6.11 Some Python Attributes Cannot Be Set on Job Objects in the JDL

Because of the upgrade to Jython 2.5, which contains a significant reworking of the Jython engine, it is no longer possible to use certain identifiers as attributes on instances of the JDL Job class. For instance,

```
class foo(Job):
    def job_started_event(self):
        self.userID = "foo"
```

results in the following job failure:

```
JobID: aspiers.jobIDtestjob.118426
Traceback (most recent call last):
  File "jobIDtestjob", line 10, in job_started_event
AttributeError: read-only attr: userID
Job 'aspiers.jobIDtestjob.118426' terminated because of failure. Reason:
AttributeError: read-only attr: userID
```

The following identifiers are known to cause problems:

- ♦ jobID
- ♦ name
- ♦ type
- ♦ userID

To work around this issue, rename any of these attributes in your JDL code.

## 7 PlateSpin Orchestrate Monitoring Issues

The following information is included in this section:

- ♦ [Section 7.1, “Some Monitored Windows Metrics Do Not Display in PlateSpin Orchestrate,” on page 10](#)

### 7.1 Some Monitored Windows Metrics Do Not Display in PlateSpin Orchestrate

The monitoring function of PlateSpin Orchestrate does not display some of the metrics of monitored Windows servers. The non-displayed metrics include:

- ♦ cpu\_idle
- ♦ cpu\_nice
- ♦ cpu\_wio
- ♦ disk\_free
- ♦ disk\_total
- ♦ load\_fifteen
- ♦ load\_five
- ♦ load\_one
- ♦ mem\_buffers
- ♦ mem\_cached
- ♦ mem\_shared
- ♦ part\_max\_used
- ♦ proc\_run

In addition to these non-displaying metrics, two others show incorrect values in resource facts:

- ♦ os\_name
- ♦ os\_release

In the Monitoring Web interface and in the VM Client, the absence of these metrics appears as a missing page. In the Development Client, the monitored resource’s fact values corresponding to these metrics all show as zero (0), except for the os\_name and os\_release metrics, which display incorrect values.

This is a known issue that exists for the Ganglia monitoring system on Windows. For more information, see the [SourceForge mailing list \(http://sourceforge.net/mailarchive/message.php?msg\\_name=B1FC9137CC38F341AC4D472659FA1376051E7ED5%40LDNPCMEU301VEUA.INTRANET.BARCAPINT.COM\)](http://sourceforge.net/mailarchive/message.php?msg_name=B1FC9137CC38F341AC4D472659FA1376051E7ED5%40LDNPCMEU301VEUA.INTRANET.BARCAPINT.COM) discussion.

## 8 PlateSpin Orchestrate Development Client Issues

The following information is included in this section:

- ◆ Section 8.1, “Using the Orchestrate Development Client in a Firewall Environment,” on page 11
- ◆ Section 8.2, “Cloning a VM onto the Datagrid Repository (zos) is Incorrectly Available as an Option in the Development Client,” on page 11
- ◆ Section 8.3, “The Development Client Displays Incorrect CPU Speed for SLES 11 SP1 Resources,” on page 12
- ◆ Section 8.4, “Maximum Instances Per VM Host Is Removed,” on page 12
- ◆ Section 8.5, “Authorization Constraint Messages Are Not Readily Visible in the Development Client,” on page 12

### 8.1 Using the Orchestrate Development Client in a Firewall Environment

Using the PlateSpin Orchestrate Development Client in a NAT-enabled firewall environment 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 NAT-enabled firewall environment, you need to use a remote desktop or VPN product.

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

### 8.2 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.

## 8.3 The Development Client Displays Incorrect CPU Speed for SLES 11 SP1 Resources

The CPU speed displayed in the Orchestrate Development Client (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 PlateSpin Orchestrate Server is also incorrect.

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.

To work around this issue, run the following command at the server command line:

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

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 following command at the bash prompt of your SLES server:

```
cat /proc/cpuinfo
```

---

## 8.4 Maximum Instances Per VM Host Is Removed

In older versions of the PlateSpin Orchestrate Server, the `resource.vm.maxinstancespervmhost` fact could be set in the Development Client, but the value was never used and so would never have any impact on server behavior. The fact has now been removed from the server fact schema and from the Development Client UI, although any non-default values set on grid resources still persist for the benefit of any custom JDL or policies that rely on them. This functionality might be fully re-implemented in the future.

## 8.5 Authorization Constraint Messages Are Not Readily Visible in the Development Client

If you unsuccessfully attempt to provision a VM whose Host/Repository selection has been designated as *Automatic*, it is possible that a policy with an authorization constraint has been associated to that VM. In this scenario, no message explaining the restriction is displayed.

To confirm that the provisioning has an authorization constraint:

- 1 In the Explorer tree of the Development client, select the VM that you want to provision.
- 2 In the Development Client toolbar, select *Provisioner* to open the provisioning monitor view for that VM.
- 3 Select the *Show Log* tab to open the provisioning log.  
Scan the log to find errors indicating that the VM could not be provisioned because of authorization constraints in its policy.

## 9 PlateSpin Orchestrate VM Client Issues

The following information is included in this section:

- ♦ [Section 9.1, “VM Does Not Start When a Local Repository Is Assigned to Multiple Hosts,” on page 13](#)
- ♦ [Section 9.2, “Not Configuring a Display Driver Triggers a Pop-Up Message,” on page 14](#)
- ♦ [Section 9.3, “Cannot Increase the Number of vCPUs on a Running Xen VM,” on page 14](#)
- ♦ [Section 9.4, “The Default Desktop Theme on SLES 10 or SLED 10 Causes a Display Problem for the VM Client,” on page 14](#)
- ♦ [Section 9.5, “Using the Orchestrate VM Client in a Firewall Environment,” on page 14](#)
- ♦ [Section 9.6, “VM Client Error Log Lists a Large Number of Display Exceptions,” on page 15](#)
- ♦ [Section 9.7, “Storage Type Options Might Not Be Visible When Modifying a Repository,” on page 15](#)
- ♦ [Section 9.8, “The Development Client Must Be Used to Set the Administrator and Domain Facts Before Cloning,” on page 15](#)

### 9.1 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 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 Development Client. 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.

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.

## 9.2 Not Configuring a Display Driver Triggers a Pop-Up Message

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.

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.

## 9.3 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.

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.

## 9.4 The Default Desktop Theme on SLES 10 or SLED 10 Causes a Display Problem for the VM Client

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. 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*.

## 9.5 Using the Orchestrate VM Client in a Firewall Environment

Using the PlateSpin Orchestrate VM Client in a NAT-enabled firewall environment is not supported for this release. The VM Development Client 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.

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

## 9.6 VM Client Error Log Lists a Large Number of Display Exceptions

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.

We are aware of the large number of exceptions occurring within this class. The errors are currently unavoidable and can be safely ignored.

## 9.7 Storage Type Options Might Not Be Visible When Modifying a Repository

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.

## 9.8 The Development Client Must Be Used to Set the Administrator and Domain Facts Before Cloning

The *Network* and *Windows* tabs have been removed from the Clone VM Wizard in the VM Client. You need to use the Development Client 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.

# 10 Virtual Machine Management General Issues

The following information is included in this section:

- ◆ [Section 10.1, “Using Autoprep When LVM Is the Volume Manager,” on page 16](#)
- ◆ [Section 10.2, “Volume Tools Hang While Scanning a Suspended Device,” on page 16](#)
- ◆ [Section 10.3, “Manually Created VM Might Display “Under Construction” on the VM Icon,” on page 16](#)
- ◆ [Section 10.4, “Canceling a VM Build Fails on a SLES 11 VM Host,” on page 16](#)
- ◆ [Section 10.5, “SUSE Linux VMs Might Attempt To Partition a Read-only Device,” on page 17](#)
- ◆ [Section 10.6, “RHEL 5 VMs Running the Kudzu Service Do Not Retain Network Interface Changes,” on page 17](#)
- ◆ [Section 10.7, “Unable to Add an Additional Disk to a Novell Cloud Manager Workload,” on page 17](#)
- ◆ [Section 10.9, “VMs Provisioned from a VM Template Are Not Restarted When a VM Host Crashes,” on page 19](#)

## 10.1 Using Autoprep When LVM Is the 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 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.

## 10.2 Volume Tools Hang While Scanning a Suspended Device

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)!
```

Because of this behavior, we recommend against using LVM and similar volume tools on a virtual machine managed by PlateSpin Orchestrate.

## 10.3 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” icon overlays the VM icon in the Orchestrate Development Client.

This flag (icon) 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.

## 10.4 Canceling a VM Build Fails on a 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 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.



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

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.

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

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.

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

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.

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.

To work around this problem, disable the Kudzu service within the RHEL VM by using the following command:

```
chkconfig --del kudzu
```

## 10.7 Unable to Add an Additional Disk to a Novell Cloud Manager Workload

Novell Cloud Manager 1.0 administrators could not add new vDisks to existing Cloud Manager workloads. However, PlateSpin Orchestrate 2.5 has been extended in [PlateSpin Orchestrate 2.5 Patch 1](http://download.novell.com/Download?buildid=JsXHTYzOwDs~) (<http://download.novell.com/Download?buildid=JsXHTYzOwDs~>) and Cloud Manager has been extended in [Novell Cloud Manager 1.0 Patch 1](http://download.novell.com/Download?buildid=JZ54WFTVEQk~) (<http://download.novell.com/Download?buildid=JZ54WFTVEQk~>) to allow this functionality.

## 10.8 Policies Applied to VM Resources Are Deleted

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”). VMs in PlateSpin Orchestrate 2.5 and later must be made standalone to receive and retain associated policies.

To work around this issue, 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>

```

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

If a VM host crashes, VMs that were provisioned from a template on that host are not restarted on another active VM host. Instead, PlateSpin Orchestrate 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, Orchestrate creates a new VM object with the `zombie_` string prepended to the VM object name.

This behavior probably occurs because the VM host crashed or the Orchestrate Agent on that host went offline while hosting a provisioned clone.

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

## 11 vSphere VM Issues in the Development Client

The following information is included in this section:

- ♦ [Section 11.1, “\(503\) Service Unavailable Errors Might Occur While Cloning vSphere VMs,” on page 19](#)
- ♦ [Section 11.2, “The Orchestrate Client Model of Multiple vCenter Datacenters Does Not Accurately Report Actual Repository Space,” on page 20](#)
- ♦ [Section 11.3, “vSphere Repository Free Space and Used Space Are Not Accurate,” on page 20](#)
- ♦ [Section 11.4, “vSphere VM Image Discovery Might Fail During Object Creation,” on page 20](#)
- ♦ [Section 11.5, “Changes to Information in vSphere Policies Are Not Enforced until the vSphere Daemon is Restarted,” on page 21](#)
- ♦ [Section 11.6, “Removing an ESX Server from the vSphere Environment Is Not Mirrored in PlateSpin Orchestrate,” on page 21](#)
- ♦ [Section 11.7, “Invalid Datastore Path Error,” on page 22](#)

### 11.1 (503) Service Unavailable Errors Might Occur While Cloning vSphere VMs

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

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

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 PlateSpin Orchestrate Agent are both installed on the same Windows Server 2003 computer.

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.

## 11.2 The Orchestrate Client Model of Multiple vCenter Datacenters Does Not Accurately Report Actual Repository Space

In a vSphere environment with multiple datacenters, if ESX hosts in separate datacenters are connected to the same shared datastore (NFS, iSCSI SAN or Fibre Channel SAN), one Orchestrate Repository object is created for each combination of datacenter and shared datastore. To illustrate:

- ♦ An ESX host named “ESX-A” resides in “Datacenter-A.” ESX-A is connected to an NFS share named “vcenterNFS.”
- ♦ An ESX host name “ESX-B” resides in “Datacenter-B.” ESX-B is connected to the same NFS share as ESX-A ( “vcenterNFS”).
- ♦ PlateSpin Orchestrate creates two Repository objects: vcenterNFS and vcenterNFS-1

Testing has shown that each of these created Orchestrate Repositories is populated with only the VMs that populated the corresponding vSphere datacenter. PlateSpin Orchestrate calculates the free and available space for a VM based only on the VMs per datacenter, rather than on the free space and available space of the shared storage where the VMs actually reside. You should be aware of this misrepresentation to avoid being misled by the displayed available options in a VM provision plan.

## 11.3 vSphere Repository Free Space and Used Space Are Not Accurate

The values for the `repository.freespace` and `repository.usedspace` facts are internally calculated by the Orchestrate server and not populated from vCenter directly. Under certain circumstances, these facts might report inaccurate values because of additional files stored on the vCenter datastore (for example,. VMs not discovered by Orchestrate, snapshot files, and so on), or datastores that are shared between multiple datacenters.

To work around this issue, you can disable the repository freespace constraint check by setting the value for the `repository.capacity` fact to “infinite” (-1).

```
<policy>
  <repository>
    <fact name="capacity"
      value="-1"
      type="Integer"
      description="Infinite repository capacity" />
  </repository>
</policy>
```

This allows Orchestrate to ignore the freespace constraint and lets vCenter later fail the provisioning adapter job if there is insufficient space available in the preferred datastore.

## 11.4 vSphere VM Image Discovery Might Fail During Object Creation

During a discover VM image operation in a vSphere environment, a race condition can occur when multiple grid objects of the same name and same type (vNICs, vDisks, vBridges) are being created simultaneously in PlateSpin Orchestrate. The name generation code tries to create a unique

Orchestrate grid name for objects that already exist (attempting to append an integer value to the end of the grid object name until it is unique in the Orchestrate grid object namespace). However, if multiple provisioning adapter discovery jobs are run concurrently, the race condition occurs: both discovery jobs pass the name generation code and one attempts to create a duplicate named grid object, evidenced in a stacktrace as follows:

```
[vsphere] Vnic list: Changed
Traceback (most recent call last):
  File "vsphere", line 4689, in handleDaemonResponse_event
  File "vsphere", line 2551, in objects_discovered_event
  File "vsphere", line 2307, in vms_discovered_event
  File "vsphere", line 2467, in update_vm_facts
  File "vsphere", line 3453, in update_vnic_facts
RuntimeError: Could not register
MBean:local:vnic=w2k3r2i586-zos107-iscsi-1_vnic1
Job 'system.vsphere.42' terminated because of failure. Reason: RuntimeError:
Could not register MBean:local:vnic=w2k3r2i586-zos107-iscsi-1_vnic1
```

If you see this traceback, we recommend that you re-run the discovery.

## 11.5 Changes to Information in vSphere Policies Are Not Enforced until the vSphere Daemon is Restarted

If you change any information in a vSphere provisioning adapter policy, such as a new or additional Web address for a vCenter server, PlateSpin Orchestrate does not recognize these changes immediately.

To work around this issue, use the Job Monitor in the Development Client to locate the *Instance Name* column of the jobs table, then find an instance named *vSphere Daemon* (or `system.vspheredaemon.xxx` in the *Job ID* column), select this job, then click the *Cancel* button at the top of the monitor page.

## 11.6 Removing an ESX Server from the vSphere Environment Is Not Mirrored in PlateSpin Orchestrate

If you use the vCenter client to remove an ESX server from the vSphere environment after PlateSpin Orchestrate has previously discovered it as a resource (VM host), Orchestrate continues to show that resource in the Explorer tree and in the `repository.vhosts` fact, making it available as a possible resource in the VM provisioning plan.

If you do not realize that the ESX resource has been removed from the vSphere environment and you run a provisioning job on a VM provisioned to the non-existent VM host, the provisioning job fails:

```
handleDaemonResponse: OpId=Ops1 OpState=Error OpError=Error looking up MOR
with
path:
Checking if networks were deleted
Checking if new networks were created
Job 'system.vsphere.1046' terminated because of failure. Reason: Error looking
up MOR with path:
```

To work around this issue, manually delete the resource and its VM host in the Development Client.

## 11.7 Invalid Datastore Path Error

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'
```

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>
```

## 12 Xen VM Issues in the Development Client

The following information is included in this section:

- ◆ [Section 12.1, “Using the XendConfig Job When Discovering a Large Number of Xen VM Hosts Continues for a Long Time,” on page 23](#)
- ◆ [Section 12.2, “Unsupported Features in the Xen Hypervisor,” on page 23](#)
- ◆ [Section 12.3, “Running xm Commands on an Old Xen VM Host Causes Server to Hang,” on page 23](#)
- ◆ [Section 12.4, “Enabling a Lock on a VM Protects only Against a Second VM from Provisioning,” on page 24](#)
- ◆ [Section 12.5, “Remote Desktop View Doesn’t Work on Fully Virtualized Linux VMs Created in SLES 11 Xen,” on page 24](#)
- ◆ [Section 12.6, “Remote Console Might Not Work On Xen Resources,” on page 24](#)
- ◆ [Section 12.7, “Remote Console to a Windows XP VM on a Xen Host Might Not Work Properly,” on page 24](#)
- ◆ [Section 12.8, “The Personalize Action Incorrectly Creates a New vDisk,” on page 25](#)
- ◆ [Section 12.9, “Deleting a vDisk from a Xen VM Does Not Delete the Disk File,” on page 25](#)
- ◆ [Section 12.10, “Moving a Xen VM with an Inaccessible Attached ISO File or Disk Creates an Unusable File,” on page 25](#)
- ◆ [Section 12.11, “Suspending a Fully Virtualized VM Makes VM Unrecoverable,” on page 25](#)
- ◆ [Section 12.12, “An Invalid Xen vNIC Model Type Might Cause Issues When a VM Is Managed in the Development Client,” on page 26](#)
- ◆ [Section 12.13, “A Suspended Xen VM Cannot Start if it Is Already Managed by a Xen Host,” on page 26](#)
- ◆ [Section 12.14, “Orchestrator Does Not Provision Xen VMs with Spaces in the VM Name,” on page 26](#)
- ◆ [Section 12.15, “A Xen VM Cannot Provision If It Uses a Physical Disk Without a Repository,” on page 26](#)

## 12.1 Using the XendConfig Job When Discovering a Large Number of Xen VM Hosts Continues for a Long Time

When the xendConfig job is used during the discovery of a very large number of Xen VM hosts (that is, Xen resources where you have installed the Orchestrate Agent), the completion of the XendConfig job can take an unnecessary amount of time to complete. This happens because, by default, an instance of the xendConfig job is started for every VM host discovered, possibly resulting in a very large number of queued xendConfig jobs.

By default, the xendConfig job is constrained to allow only one instance of the job to run at a time, causing all the other xendConfig job instances to queue.

The following constraint from the `xendConfig.policy` file causes all the XendConfig job instances to run one at a time, rather than concurrently.

```
<constraint type="start" >
  <lt fact="job.instances.active"
      value="1"
      reason="Only 1 instance of this job can run at a time" />
</constraint>
```

If you need to work around this issue to accommodate a large Xen environment, you can temporarily remove or comment out this constraint from the xendConfig policy, but you must ensure that no other Orchestrate administrator runs this job at the same time. Doing so might result in corruption of the `/etc/xen/xend-config.sxp` file because two xendConfig job instances could attempt to concurrently modify this config file.

## 12.2 Unsupported Features in the Xen Hypervisor

The checkpoint and restore features on the xen provisioning adapter only suspend and resume the specified VM. Xen does not support taking normal snapshots as other hypervisors do.

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

The Xen provisioning adapter uses `xm` commands to perform basic VM 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.

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
```

## 12.4 Enabling a 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, then that node loses network connectivity to the Orchestrate Server, a reprovisioning of the VM fails because the lock is 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.

## 12.5 Remote Desktop View Doesn't Work on Fully Virtualized Linux VMs Created in SLES 11 Xen

If you try to connect to a fully virtualized Linux VM by using the Development Client, VM Client, or any other utility that uses `vncviewer.jar`, the remote view is garbled or does not stay connected.

To work around the problem, use a regular VNC client or the Xen `virt-viewer`.

## 12.6 Remote Console Might Not Work On Xen Resources

The VNC client you launch from Novell Cloud Manager or from the PlateSpin Orchestrate Development Client does not work if the host name of the Xen host (that is, the `resource.vnc.ip` fact) is not resolvable to an IP address.

You can determine the current host name with the `hostname --fqdn` command at the bash prompt of the host.

If the host name is unresolvable, PlateSpin Orchestrate might try to supply an alternate IP address or DNS name. You can edit the `/etc/hosts` file to ensure that the DNS name points to the actual IP address.

## 12.7 Remote Console to a Windows XP VM on a Xen Host Might Not Work Properly

The VNC client (`vncviewer.jar`) you launch from the PlateSpin Orchestrate Development Client to connect to a Windows XP VM running on a Xen Host can occasionally render a garbled desktop UI.



To work around the problem, update the jar file in `/opt/novell/zenworks/zos/clients/lib/gpl/vncviewer.jar` with the jar file available at <http://www.tightvnc.com/ssh-java-vnc-viewer.php> (<http://www.tightvnc.com/ssh-java-vnc-viewer.php>).

## 12.8 The Personalize Action Incorrectly Creates a New vDisk

If you create a vDisk on a Xen VM or VM template and then you execute the *Personalize* action, the configuration for the new vDisk is saved in `config.xen` and a new disk image is created on the Xen host. Normally, only the *Save Config* action should be allowed to create a new disk object. Although *Save Config* still works in this scenario, *Personalize* should not. The issue will be addressed in the next release.

## 12.9 Deleting a vDisk from a Xen VM Does Not Delete the Disk File

If you use the Orchestrate Client to delete a vDisk from a Xen VM that has several vDisk images attached, then use the *Save Config* action to save the deletion, the vDisk is removed from the Explorer tree and from the Xen config file, but the disk image is not deleted.

If you want the disk image to be deleted, you must do so manually (that is, outside PlateSpin Orchestrate) from the file system or storage container where the image is located. For Xen, you can do this by using standard Linux file operations. For other hypervisors, you can do this by using the hypervisor's management interface.

## 12.10 Moving a Xen VM with an Inaccessible Attached ISO File or Disk Creates an Unusable File

If you move a Xen VM that has an attached ISO file whose location is inaccessible or unknown to PlateSpin Orchestrate, Orchestrate creates a file that takes the place of the ISO, but the file is not the actual ISO. The same thing occurs if you attach a disk file located in an undefined repository.

Before you use Orchestrate to attempt moving the VM disks, we recommend that you remove any VM's ISO disk that does not reside in the same repository.

## 12.11 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 drivers 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
```

## 12.12 An Invalid Xen vNIC Model Type Might Cause Issues When a VM Is Managed in the Development Client

Although restriction of valid vNIC types for Xen VMs occurs in the Orchestrate VM Client, the Development Client allows editing of the type (in the *Constraints* table under the *Constraints/Facts* tab of the Admin view) to any string. The Development Client accepts any string as a valid vNIC type, even if it is not supported by the VM Client. In this situation, the VM can be provisioned, but it does so in an unstable state, such as running indefinitely after being provisioned or being unable to launch a remote session to the VM from the Development Client.

To work around this situation, you can manually shut down or remove the VM by using the `xm` command on the host where it was provisioned.

## 12.13 A Suspended Xen VM Cannot Start if it Is Already Managed by a Xen Host

If PlateSpin Orchestrate discovers a VM in a suspended state (that is, a checkpoint file exists for it) on the Xen VM host, Orchestrate cannot start and provision that VM.

To work around the issue, run the `xm delete` command at the Xen host to remove the VM from management by the Xen host. The VM then becomes manageable by PlateSpin Orchestrate.

## 12.14 Orchestrator Does Not Provision Xen VMs with Spaces in the VM Name

Even though the Xen hypervisor lets you create a VM with spaces in its name and PlateSpin Orchestrate successfully discovers such a named VM image (along with the similarly named VM directory, and VM config file), provisioning a VM with spaces in its name from PlateSpin Orchestrate or from the Orchestrate Server command line fails. Specifically, the failure occurs when the `xm` command runs.

We recommend that you rename all such VMs before provisioning so that no spaces exist in the name.

## 12.15 A Xen VM Cannot Provision If It Uses a Physical Disk Without a Repository

If the xen provisioning adapter attempts to discover a VM that has an attached physical type vDisk (that is, a vDisk defined in the `config.xen` file with a `phy://` location), the provisioning adapter fails to discover the physical disk device.

A patch has been provided to address this problem. Download and apply [PlateSpin Orchestrate 2.5 Patch 1](http://www.novell.com/support/php/search.do?cmd=displayKC&docType=kc&externalId=InfoDocument-patchbuilder-readme5080146&sliceId=&docTypeID=DT_SUSESDB_PSDB_1_1&dialogID=87082973&stateId=0%20%20174255354) ([http://www.novell.com/support/php/search.do?cmd=displayKC&docType=kc&externalId=InfoDocument-patchbuilder-readme5080146&sliceId=&docTypeID=DT\\_SUSESDB\\_PSDB\\_1\\_1&dialogID=87082973&stateId=0%20%20174255354](http://www.novell.com/support/php/search.do?cmd=displayKC&docType=kc&externalId=InfoDocument-patchbuilder-readme5080146&sliceId=&docTypeID=DT_SUSESDB_PSDB_1_1&dialogID=87082973&stateId=0%20%20174255354)), available from Novell Support.

## 13 Hyper-V VM Issues in the Development Client

The following information is included in this section:

- ♦ [Section 13.1, “The Admin Password Fact Must Be Set on a Hyper-V Windows VM for Sysprep to Work,”](#) on page 27
- ♦ [Section 13.2, “The VNC Console Does Not Work on a Hyper-V VM when Invoked from Cloud Manager,”](#) on page 27
- ♦ [Section 13.3, “Remote Console Works Intermittently on a Hyper-V VM Host,”](#) on page 28
- ♦ [Section 13.4, “Hyper-V Provisioning Jobs Fail When Several Jobs Are Started Simultaneously,”](#) on page 28
- ♦ [Section 13.5, “Limitations of Linux VMs as Guests on Hyper-V,”](#) on page 28
- ♦ [Section 13.6, “The Development Client Might Report an Incorrect Capacity for a Cluster Shared Volume in a Hyper-V Environment,”](#) on page 28

Other ongoing issues for Hyper-V VMs are documented in [“Configuring the hyperv Provisioning Adapter and Hyper-V VMs”](#) in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

### 13.1 The Admin Password Fact Must Be Set on a Hyper-V Windows VM for Sysprep to Work

As with other VMs provisioned by PlateSpin Orchestrate, 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 [“Autoprep and Sysprep”](#) section of the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

### 13.2 The VNC Console Does Not Work on a Hyper-V VM when Invoked from Cloud Manager

If you invoke the VNC console for a Hyper-V VM (referred to as a “workload”) from Novell Cloud Manager, the VNC console does not launch.

Installing the PlateSpin Orchestrate 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 client to the VM:

- 1 From Explorer tree in the Development Client, 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*.

## 13.3 Remote Console Works Intermittently on a Hyper-V VM Host

Testing has shown that launching and using a remote console VNC session on a Hyper-V VM host from Novell Cloud Manager sometimes fails.

We recommend that you use the latest release of any VNC server software available. If the problem persists, close the remote console window and try relaunching the remote session.

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

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.

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 Orchestrate 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 *PlateSpin Orchestrate 2.5 Development Client Reference*.

## 13.5 Limitations of Linux VMs as Guests on Hyper-V

PlateSpin Orchestrate does not support the *Create Template* or *Clone* actions for Linux-based Hyper-V VMs.

## 13.6 The Development Client Might Report an Incorrect Capacity for a Cluster Shared Volume in a Hyper-V Environment

If you perform the *Discover VM Hosts and Repositories* action in a Hyper-V environment with a Cluster Shared Volume (CSV), PlateSpin Orchestrate might occasionally report the `repository.capacity` value and the `repository.freespace` value as 0 MB. If this occurs, this repository is not available to perform provisioning operations (such as creating a template) in Orchestrate.

To work around the issue, change the `repository.capacity` value either to -1 (infinite) or to the actual capacity of the CSV.

# 14 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 © 2008-2010 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.

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.