

# PlateSpin Orchestrate 2.6 Readme

December 8, 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 2.6 was initially released.

- ◆ [Section 1, “Network File System Issues,” on page 1](#)
- ◆ [Section 2, “YaST Issues,” on page 2](#)
- ◆ [Section 3, “Installation Issues,” on page 2](#)
- ◆ [Section 4, “Upgrade Issues,” on page 3](#)
- ◆ [Section 5, “PlateSpin Orchestrate Server Issues,” on page 4](#)
- ◆ [Section 6, “PlateSpin Orchestrate Monitoring Issues,” on page 6](#)
- ◆ [Section 7, “PlateSpin Orchestrate Development Client Issues,” on page 7](#)
- ◆ [Section 8, “PlateSpin Orchestrate VM Client Issues,” on page 8](#)
- ◆ [Section 9, “Virtual Machine Management General Issues,” on page 10](#)
- ◆ [Section 10, “vSphere VM Issues in the Development Client,” on page 11](#)
- ◆ [Section 11, “Xen VM Issues in the Development Client,” on page 13](#)
- ◆ [Section 12, “Hyper-V VM Issues in the Development Client,” on page 15](#)
- ◆ [Section 13, “Legal Notices,” on page 16](#)

## 1 Network File System Issues

The following information is included in this section:

- ◆ [Section 1.1, “Orchestrate Agent Fails to Set the UID on Files Copied from the Datagrid,” on page 1](#)

### 1.1 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.

---

## 2 YaST Issues

The following information is included in this section:

- ♦ [Section 2.1, “YaST Uninstall Feature Is Not Supported,” on page 2](#)

### 2.1 YaST Uninstall Feature Is Not Supported

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

## 3 Installation Issues

The following information about installation is included in this section:

- ♦ [Section 3.1, “Configuring the Orchestrate Agent on RHEL 4 Machines Fails,” on page 2](#)
- ♦ [Section 3.2, “The Orchestrate Agent Should Not Be Installed on a vCenter Server,” on page 3](#)

### 3.1 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 2.6 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/).

## 3.2 The Orchestrate Agent Should Not Be Installed on a vCenter Server

We recommend that you do not install the Orchestrate Agent on a machine where the vCenter server is running because the agent's potential processing needs might put an unsatisfactory load on the vCenter server. Testing has also shown that a race condition can result if the Orchestrate Agent and the vCenter server are started simultaneously, possibly resulting in the failure of the vCenter server to start properly.

You can work around this issue by installing the agent on another machine that has a remote connection to the vCenter server. For more information see [“Discovering Enterprise Resources in Multiple vSphere Environments”](#) in the *PlateSpin Orchestrate 2.6 Virtual Machine Management Guide*.

## 4 Upgrade Issues

The following information is included in this section:

- ◆ [Section 4.1, “Audit Database Values Are Not Preserved in an Upgrade,” on page 3](#)
- ◆ [Section 4.2, “A Clone Does Not Inherit the Policy Associations of Its Upgraded Parent VM Template,” on page 3](#)
- ◆ [Section 4.3, “A Xen VM Template Has an Extra Clone after Upgrade,” on page 4](#)
- ◆ [Section 4.4, “A Clone Provisioned from a Xen VM Template Loses Template Dependency on Upgrade,” on page 4](#)

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

If you upgrade the PlateSpin Orchestrate Server to version 2.6, 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.

### 4.2 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.

### 4.3 A Xen VM Template Has an Extra Clone after Upgrade

If your Orchestrate 2.5 environment includes a provisioned Xen VM template you will notice that an extra VM clone of that template is created after an upgrade to Orchestrate 2.6. This clone is associated with the same configuration and vDisks as the template.

Because this “invalid” clone is unexpected, you might be inclined to delete it. If you do so, however, the files for the original template are deleted, making it impossible to create more clones.

### 4.4 A Clone Provisioned from a Xen VM Template Loses Template Dependency on Upgrade

If you provision a clone from a Xen VM template created in PlateSpin Orchestrate 2.5 and then you subsequently upgrade Orchestrate to version 2.6, the clone loses its dependency on the VM template and disappears from the Development Client.

If you run a VM discovery after the upgrade, the VM instance reappears, but is marked as a “zombie” VM without any association to the template VM.

To work around this issue, any VMs that are provisioned directly from a template (where `resource.provision.automatic = True`), must either be disassociated from the template or destroyed.

## 5 PlateSpin Orchestrate Server Issues

The following information is included in this section:

- ♦ [Section 5.1, “The Orchestrate Server Might Require Caching of Computed Facts in a Grid with Large Numbers of Resources,” on page 4](#)
- ♦ [Section 5.2, “Calling terminate\(\) from within a Job Class Allows the JDL Thread Execution to Continue,” on page 5](#)
- ♦ [Section 5.3, “Deploying Components Might Fail Intermittently,” on page 5](#)
- ♦ [Section 5.4, “Some Python Attributes Cannot Be Set on Job Objects in the JDL,” on page 6](#)
- ♦ [Section 5.5, “Java Programs Run with the JDL exec Feature Might Hang,” on page 6](#)

### 5.1 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.

---

## 5.2 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`.

## 5.3 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.

## 5.4 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.

## 5.5 Java Programs Run with the JDL exec Feature Might Hang

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

To work around this issue, turn off the enhanced ExecWrapper, as follows:

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

---

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

---

For more information about the JDL Exec class, see the [Orchestrate 2.6 JDL documentation \(http://www.novell.com/documentation/cloudmanager1/resources/javadoc/com/novell/zos/jdl/Exec.html\)](http://www.novell.com/documentation/cloudmanager1/resources/javadoc/com/novell/zos/jdl/Exec.html).

## 6 PlateSpin Orchestrate Monitoring Issues

The following information is included in this section:

- ♦ [Section 6.1, “Monitoring for RHEL 4, RHEL 5, and SLES 9 Resources Is Not Included in the Installation Packages,” on page 7](#)

## 6.1 Monitoring for RHEL 4, RHEL 5, and SLES 9 Resources Is Not Included in the Installation Packages

The PlateSpin Orchestrate 2.6 installation media does not include the RHEL4, RHEL5, or SLES 9 monitoring packages.

If you want to monitor RHEL or SLES 9 resources, we recommend that you download Ganglia 3.1.7 from the [SourceForge \(http://sourceforge.net/projects/ganglia/files/ganglia%20monitoring%20core/3.1.7/\)](http://sourceforge.net/projects/ganglia/files/ganglia%20monitoring%20core/3.1.7/) Web site and install it on the resources to be monitored. Create a `.conf` file similar to one that exists on a SLES machine, editing the node name in the file so that the monitoring metrics display for the resource in the Orchestrate Development Client.

## 7 PlateSpin Orchestrate Development Client Issues

The following information is included in this section:

- ♦ [Section 7.1, “Cloning a VM onto the Datagrid Repository \(zos\) is Incorrectly Available as an Option in the Development Client,” on page 7](#)
- ♦ [Section 7.2, “Maximum Instances Per VM Host Is Removed,” on page 7](#)
- ♦ [Section 7.3, “Memory Changes in the VM Host Might Not Be Accurately Displayed,” on page 7](#)

### 7.1 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.

### 7.2 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 implemented in the future.

### 7.3 Memory Changes in the VM Host Might Not Be Accurately Displayed

The memory available for provisioning operations is exposed by the `vmhost.memory.available` fact. You can see this value in the VM host admin view of the Development Client under the *Guest VM Monitor Information* section of the *Info* page. You can also see the value on the *Constraints/Facts* page of the Development Client.

In determining this value, the Orchestrate Server does not account for any interim change to the physical RAM available on a VM host until you perform a rediscovery of the VM host(s) or manually update the `vmhost.memory.max` fact.

## 8 PlateSpin Orchestrate VM Client Issues

The following information is included in this section:

- ♦ [Section 8.1, “VM Does Not Start When a Local Repository Is Assigned to Multiple Hosts,” on page 8](#)
- ♦ [Section 8.2, “Not Configuring a Display Driver Triggers a Pop-Up Message,” on page 8](#)
- ♦ [Section 8.3, “Cannot Increase the Number of vCPUs on a Running Xen VM,” on page 9](#)
- ♦ [Section 8.4, “The Default Desktop Theme on SLES 10 or SLED 10 Causes a Display Problem for the VM Client,” on page 9](#)
- ♦ [Section 8.5, “Using the Orchestrate VM Client in a Firewall Environment,” on page 9](#)
- ♦ [Section 8.6, “VM Client Error Log Lists a Large Number of Display Exceptions,” on page 10](#)
- ♦ [Section 8.7, “Storage Type Options Might Not Be Visible When Modifying a Repository,” on page 10](#)
- ♦ [Section 8.8, “The Development Client Must Be Used to Set the Administrator and Domain Facts Before Cloning,” on page 10](#)
- ♦ [Section 8.9, “Xen VMs Created in the VM Client Are Not Created in the Correct Location,” on page 10](#)

### 8.1 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.

### 8.2 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.



### 8.3 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.

### 8.4 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*.

### 8.5 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.

### 8.6 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.

## 8.7 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.

## 8.8 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.

## 8.9 Xen VMs Created in the VM Client Are Not Created in the Correct Location

When you use the PlateSpin Orchestrate VM Client to create a new Xen VM, the client places that VM at the root of the filesystem instead of the selected repository path.

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

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

# 9 Virtual Machine Management General Issues

The following information is included in this section:

- ♦ [Section 9.1, “Using Autoprep When LVM Is the Volume Manager,” on page 11](#)
- ♦ [Section 9.2, “Canceling a VM Build Fails on a SLES 11 VM Host,” on page 11](#)

## 9.1 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.

## 9.2 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 vSphere VM Issues in the Development Client

The following information is included in this section:

- ♦ [Section 10.1, “The Orchestrate Client Model of Multiple vCenter Datacenters Does Not Accurately Report Actual Repository Space,” on page 11](#)
- ♦ [Section 10.2, “vSphere Repository Free Space and Used Space Are Not Accurate,” on page 12](#)
- ♦ [Section 10.3, “vSphere VM Image Discovery Might Fail During Object Creation,” on page 12](#)
- ♦ [Section 10.4, “Changes to Information in vSphere Policies Are Not Enforced until the vSphere Daemon is Restarted,” on page 13](#)

## 10.1 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.

## 10.2 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.

## 10.3 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.

## 10.4 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 Xen VM Issues in the Development Client

The following information is included in this section:

- ♦ [Section 11.1, “Enabling a Lock on a VM Protects only Against a Second VM from Provisioning,” on page 13](#)
- ♦ [Section 11.2, “Remote Desktop View Doesn’t Work on Fully Virtualized Linux VMs Created in SLES 11 Xen,” on page 14](#)
- ♦ [Section 11.3, “Remote Console to a Windows XP VM on a Xen Host Might Not Work Properly,” on page 14](#)
- ♦ [Section 11.4, “Deleting a vDisk from a Xen VM Does Not Delete the Disk File,” on page 14](#)
- ♦ [Section 11.5, “Moving a Xen VM with an Inaccessible Attached ISO File or Disk Creates an Unusable File,” on page 14](#)
- ♦ [Section 11.6, “An Invalid Xen vNIC Model Type Might Cause Issues When a VM Is Managed in the Development Client,” on page 14](#)

### 11.1 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.

## 11.2 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`.

## 11.3 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>).

## 11.4 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.

## 11.5 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.

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

The following information is included in this section:

- ♦ [Section 12.1, “Remote Console Works Intermittently on a Hyper-V VM Host,” on page 15](#)
- ♦ [Section 12.2, “Hyper-V Provisioning Jobs Fail When Several Jobs Are Started Simultaneously,” on page 15](#)
- ♦ [Section 12.3, “Limitations of Linux VMs as Guests on Hyper-V,” on page 15](#)
- ♦ [Section 12.4, “Some Discovered Hyper-V VMs Might Have vNICs that Cannot Be Edited in the Development Client,” on page 16](#)

Other ongoing issues for Hyper-V VMs are documented in “[Configuring the hyperv Provisioning Adapter and Hyper-V VMs](#)” in the *PlateSpin Orchestrate 2.6 Virtual Machine Management Guide*.

### 12.1 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.

### 12.2 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 2.6 Development Client Reference*.

### 12.3 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.

## 12.4 Some Discovered Hyper-V VMs Might Have vNICs that Cannot Be Edited in the Development Client

PlateSpin Orchestrate might discover some Hyper-V VMs whose `vnic.type` fact value is `EmulatedEthernetPort`. In Hyper-V, this is known as a *Legacy Network Adapter*, and it must be specifically added to a VM using the Hyper-V Settings Wizard.

You cannot modify the property settings (such as *MAC Address*, *vBridge*, *Network*, and so on) of a legacy vNIC in the Development Client unless that vNIC has been previously connected to a network in Hyper-V Manager, where the proper string identifiers are associated with that vNIC.

When Orchestrate discovers a legacy vNIC without a network connection, its `vnic.connection` fact is set to null, which can result in the failure of a business service request in Novell Cloud Manager. If this happens, make sure that the vNIC has a proper network connection in the Hyper-V Manager, then rediscover the associated VM using the Development Client.

## 13 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.