

## VM Client Guide and Reference

# Novell® PlateSpin® Orchestrator

**2.0**

February 20, 2009

[www.novell.com](http://www.novell.com)



## Legal Notices

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

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

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

Copyright © 2007-2009 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed on the [Novell Legal Patents Web page \(http://www.novell.com/company/legal/patents/\)](http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.  
404 Wyman Street, Suite 500  
Waltham, MA 02451  
U.S.A.  
[www.novell.com](http://www.novell.com)

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

## **Novell Trademarks**

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

## **Third-Party Materials**

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



# Contents

<b>About This Guide</b>	<b>9</b>
<b>1 Overview</b>	<b>11</b>
1.1 What Is the PlateSpin Orchestrate VM Client? . . . . .	11
1.2 Understanding the Virtual Machine Life Cycle . . . . .	12
1.3 Using This Guide to Manage VMs . . . . .	13
<b>2 Understanding the VM Client Interface</b>	<b>15</b>
2.1 Starting the VM Client Interface . . . . .	15
2.1.1 Starting the Interface on a Windows Workstation . . . . .	15
2.1.2 Starting the Interface on a Linux Workstation . . . . .	17
2.2 Navigating the Welcome Page . . . . .	18
2.2.1 Welcome Page Main Menu Options . . . . .	19
2.2.2 Welcome Page Buttons . . . . .	21
2.3 Navigating the Inventory Views . . . . .	23
2.3.1 Inventory Lists . . . . .	23
2.3.2 Inventory Details . . . . .	24
2.4 Saving Changes in the VM Client . . . . .	25
2.5 Understanding the VM Client View and Details Editors . . . . .	25
<b>3 Creating and Setting Up Virtual Machines</b>	<b>27</b>
3.1 Registering VM Hosts . . . . .	27
3.2 Discovering Registered VM Hosts . . . . .	29
3.3 Discovering VMs of the Registered VM Hosts . . . . .	29
3.4 Registering VMs . . . . .	30
3.5 Creating a Xen VM . . . . .	31
3.6 Installing a VM . . . . .	38
3.7 Installing the PlateSpin Orchestrate Agent on a VM . . . . .	38
3.8 Installing the Monitoring Agent on a VM . . . . .	39
3.9 Creating a Template from a VM . . . . .	40
3.10 Adding Repositories . . . . .	41
3.11 Associating Repositories . . . . .	43
3.12 Removing Repositories . . . . .	43
3.12.1 Removing a Repository from a Host Server . . . . .	44
3.12.2 Removing a Repository from the VM Client . . . . .	44
<b>4 Configuring Virtual Machines</b>	<b>45</b>
4.1 Editing VMs or Templates . . . . .	45
4.2 Moving VMs . . . . .	53
4.2.1 Prerequisites . . . . .	53
4.2.2 Moving a VM . . . . .	53
4.3 Migrating VMs . . . . .	54
4.3.1 Prerequisites . . . . .	54
4.3.2 Migrating a VM . . . . .	55
4.4 Deleting VMs . . . . .	56

4.5	Deleting Templates . . . . .	57
4.6	Cloning Templates . . . . .	57
4.7	Detaching Clones from Templates . . . . .	58
4.8	Creating and Populating Groups . . . . .	59
4.8.1	Understanding Groups . . . . .	59
4.8.2	Creating a Group . . . . .	60
4.8.3	Adding Members to an Existing Group . . . . .	60
4.8.4	Removing Members from a Group . . . . .	61
4.8.5	Deleting a Group . . . . .	61
<b>5</b>	<b>Managing Virtual Machines</b>	<b>63</b>
5.1	Logging In to a Virtualization Grid . . . . .	63
5.2	Understanding a VM's Status . . . . .	64
5.3	Canceling a VM's In-Progress Action . . . . .	65
5.4	Using Feedback to Manage VMs . . . . .	65
5.5	Using Feedback to Manage Host Servers . . . . .	67
5.6	Starting a VM . . . . .	67
5.7	Shutting Down a VM . . . . .	69
5.8	Accessing a VM's Console . . . . .	70
5.9	Pausing a VM . . . . .	70
5.10	Resuming a Paused VM . . . . .	70
5.11	Suspending a VM . . . . .	71
5.12	Resuming a Suspended VM . . . . .	71
5.13	Resynchronizing a Host Server's State . . . . .	72
5.14	Resynchronizing a VM's State . . . . .	72
5.15	Using Groups to Manage Your VMs . . . . .	73
5.16	Using the Error Log View . . . . .	73
5.17	Using the Progress View . . . . .	73
5.18	Resetting the VM Client Perspectives . . . . .	74
5.19	Viewing and Editing VM Client Details . . . . .	74
5.19.1	Virtual Machine Details . . . . .	75
5.19.2	Host Server Details . . . . .	75
5.19.3	Storage Details . . . . .	76
5.19.4	Template Details . . . . .	76
5.19.5	Details Tabs . . . . .	77
<b>6</b>	<b>Troubleshooting Virtual Machine Management</b>	<b>89</b>
6.1	Adding Multiple NPIV Disks to a VM . . . . .	89
6.2	Accessing VMs and Host Servers . . . . .	89
6.3	Cloning VMs . . . . .	90
6.4	Moving or Migrating the VMs . . . . .	90
6.5	Registering VM Hosts . . . . .	91
<b>A</b>	<b>VM Installation Sources</b>	<b>93</b>
A.1	Disk Installation Sources . . . . .	93
A.2	Guest Operating System Installation Sources . . . . .	93
A.2.1	SUSE . . . . .	93
A.2.2	NetWare . . . . .	94
A.2.3	Other Linux Installation Sources . . . . .	94

<b>B NPIV Terminology</b>	<b>95</b>
<b>C Index</b>	<b>97</b>



# About This Guide

This guide introduces the PlateSpin® Orchestrate VM Client, including its basic administration environment, which is accessed through an Eclipse\* rich client platform. This guide provides an introductory overview of the VM Client, and explains how to install, monitor, and manage VMs.

The guide is organized as follows:

- ♦ Chapter 1, “Overview,” on page 11
- ♦ Chapter 2, “Understanding the VM Client Interface,” on page 15
- ♦ Chapter 3, “Creating and Setting Up Virtual Machines,” on page 27
- ♦ Chapter 4, “Configuring Virtual Machines,” on page 45
- ♦ Chapter 5, “Managing Virtual Machines,” on page 63
- ♦ Chapter 6, “Troubleshooting Virtual Machine Management,” on page 89
- ♦ Appendix A, “VM Installation Sources,” on page 93
- ♦ Appendix B, “NPIV Terminology,” on page 95
- ♦ Appendix C, “Index,” on page 97

## Audience

This book is for data center administrators. It assumes that users of the product have the following background:

- ♦ General understanding of network operating environments and systems architecture
- ♦ Knowledge of basic Linux\* shell commands, the Windows\* command prompt, and text editors

## Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to [www.novell.com/documentation/feedback.html](http://www.novell.com/documentation/feedback.html) (<http://www.novell.com/documentation/feedback.html>) and enter your comments there.

## Additional Documentation

In addition to this *Administration Guide*, PlateSpin Orchestrate 2.0 includes the following additional guides that contain valuable information about the product:

- ♦ *PlateSpin Orchestrate 2.0 Getting Started Reference*
- ♦ *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*
- ♦ *PlateSpin Orchestrate 2.0 Upgrade Guide*
- ♦ *PlateSpin Orchestrate 2.0 High Availability Configuration Guide*
- ♦ *PlateSpin Orchestrate 2.0 Administrator Reference*
- ♦ *PlateSpin Orchestrate 2.0 Command Line Reference*
- ♦ *PlateSpin Orchestrate 2.0 Virtual Machine Management Guide*

- ◆ *PlateSpin Orchestrate 2.0 Development Client Reference*
- ◆ *PlateSpin Orchestrate 2.0 Developer Guide and Reference*
- ◆ *PlateSpin Orchestrate 2.0 Job Manager Guide*

## **Documentation Updates**

For the most recent version of this guide, visit the [PlateSpin Orchestrate 2.0 documentation Web site](http://www.novell.com/documentation/ps_orchestrate20/) ([http://www.novell.com/documentation/ps\\_orchestrate20/](http://www.novell.com/documentation/ps_orchestrate20/)).

## **Documentation Conventions**

In Novell<sup>®</sup> documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (<sup>®</sup>, <sup>™</sup>, etc.) denotes a Novell trademark. An asterisk (\*) denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux, should use forward slashes as required by your software.

Review the following sections to understand the PlateSpin® Orchestrate VM Client:

- ♦ [Section 1.1, “What Is the PlateSpin Orchestrate VM Client?,” on page 11](#)
- ♦ [Section 1.2, “Understanding the Virtual Machine Life Cycle,” on page 12](#)
- ♦ [Section 1.3, “Using This Guide to Manage VMs,” on page 13](#)

## 1.1 What Is the PlateSpin Orchestrate VM Client?

The PlateSpin Orchestrate VM Client is a management interface that administrators and architects can use to manage the life cycle of the virtual machines (VMs) in your enterprise, including creating, starting, stopping, migrating, and deleting VMs. PlateSpin Orchestrate lets you better align IT to your business, control costs, and minimize risks across all VM platforms in the data center. You can increase the functionality of your data center by fully leveraging VMs as a usable resource.

The VM Client provides management of VMs from several different virtualization host types, including SUSE® Linux Enterprise Server (SLES) Xen\*, VMware\* ESX Server, and Microsoft\* Windows Server\* 2008 with Hyper-V\*.

The VM Client also allows you to manage both VMs and VM hosts in your data center. A VM host is defined as a machine running a [supported hypervisor agent](#).

The VM Client provides the following:

- ♦ A VM creation wizard that covers a VM’s installation source and mode, its virtualization mode, hardware configurations, storage devices, and network configuration
- ♦ A VM editing wizard, which can also be used to edit existing VMs, even those that are in a running state
- ♦ The ability to discover existing VMs residing on your VM host servers
- ♦ The ability to discover VM host servers in your virtualization grid; a machine can be a host server if it is running hypervisor software
- ♦ Methods for controlling VMs, such as starting, stopping, pausing, and suspending
- ♦ The ability to migrate a running VM from one host server to another in real time
- ♦ The ability to install and manage the PlateSpin Orchestrate Agent on your VMs
- ♦ The ability to make templates of VMs, including making clones of the templates
- ♦ The ability to move a VM’s repository from one host server to another
- ♦ The ability to group VMs, VM hosts, repositories, and templates for easier management
- ♦ Detailed views of the properties for each VM, host server, repository, and template
- ♦ The ability to monitor your VMs and host servers in real time
- ♦ Event logs for VMs, host servers, and templates

The VM Client enhances the functionality of the PlateSpin Orchestrate Server, enabling it to control VMs in your data center.

PlateSpin Orchestrate also provides the Development Client, where you can manage and control **data center jobs** and processes through the application of rules, policies, scheduling, utilization, and billing data center resources. For more information on the Development Client, see the *PlateSpin Orchestrate 2.0 Development Client Reference*.

The VM Client and the Development Client management interfaces work together to help you to maximize the use of VMs in your data center. You can have fewer physical machines while giving your data center many additional resources, and you can manage the physical VM host machines and their VMs in your data center.

The Orchestrate Server manages resources to perform work. It does this through automated jobs (written in Jython) that in turn are broken down into joblets that are distributed among multiple resources. For more information about PlateSpin Orchestrate and Orchestrate Server operations, see “**What You Should Know**” in the *PlateSpin Orchestrate 2.0 Developer Guide and Reference*.

In addition, as calls for resources lessen and resources are released, the Orchestrate Server evaluates the capabilities of the active resources (VM or physical machine) and keeps the best resource for the job. This might include any of the currently running VMs.

## 1.2 Understanding the Virtual Machine Life Cycle

The life cycle of a VM includes its creation, testing, modifications, use in your environment, and removal when it's no longer needed.

For example, in setting up your VM environment, you might want to first create basic VMs from which you can create templates. Then, to enable the most efficient use of your current hardware capabilities, you can use those templates to create the many different specialized VMs that you need to perform the various jobs. You can create and manage VM-related jobs through the Development Client interface.

Life cycle functions are performed one at a time per given VM as jobs on the host server in order to prevent conflicts in using the VM. Life cycle events include:

- ◆ Creating a VM
- ◆ Starting and stopping a VM
- ◆ Pausing, suspending, and resuming a VM
- ◆ Installing the Orchestrate Agent on a VM
- ◆ Creating a template from a VM
- ◆ Using the VM (starting, stopping, pausing, suspending, restarting, and shutting down)
- ◆ Running jobs for the VM
- ◆ Editing a VM
- ◆ Editing a template
- ◆ Moving a stopped VM to another host server
- ◆ Migrating a running VM to another host server
- ◆ Resynchronizing a VM to ensure that the state of the VM displayed in the Orchestrate Development Client is accurate
- ◆ Cloning a VM from a template

## 1.3 Using This Guide to Manage VMs

After installing the PlateSpin Orchestrate VM Client, do the following:

1. Become familiar with the VM Client interface.
2. Start the VM Client interface.
3. Register existing VM host servers.
4. Discover the registered host servers.
5. Discover the registered VMs.
6. Log in to a virtualization grid.
7. Create your VMs.
8. Install the VMs.
9. Install the Orchestrate Agent on the VMs.
10. Create templates of the VMs.
11. Add repositories.
12. Configure the VMs:
  - ♦ Edit a VM or template.
  - ♦ Delete a VM.
  - ♦ Delete a template.
  - ♦ Clone a template.
  - ♦ Move a VM.
  - ♦ Migrate a VM.
  - ♦ Create groups for managing VMs, host servers, repositories, and templates.
13. Use the VMs:
  - ♦ Understand the various VM statuses.
  - ♦ Start VMs.
  - ♦ Stop VMs.
  - ♦ View a VM's server console.
  - ♦ Pause a VM.
  - ♦ Suspend a VM.
  - ♦ Resynchronize a host server with the Development Client.
  - ♦ Resynchronize a VM with the Development Client.
  - ♦ View logging details that are fed back from the Orchestrate Server.
  - ♦ View the error log for a VM.
  - ♦ Open the Progress view for a VM while a life cycle function is running.
  - ♦ View a VM's details.
  - ♦ View a host server's details.
  - ♦ View a repository's details.
  - ♦ View a template's details.

14. [Troubleshoot VMs.](#)
15. [Find operating system installation sources for VMs.](#)

# Understanding the VM Client Interface

# 2

The PlateSpin® Orchestrate VM Client interface, built on an Eclipse-based rich client platform, uses a graphical user interface (GUI) to help you create, store, edit, and use VMs. This interface can be installed on both Windows and Linux administration devices.

Multiple administrators can each be running an instance of the VM Client that is logged in to the same VM host. In other words, multiple administrators can simultaneously manage the VMs on a particular host server. When any administrator performs an action against a VM Client object, all VM Clients that are opened to that host server are updated accordingly.

Review the following sections for tips on navigating and using the VM Client interface:

- ♦ [Section 2.1, “Starting the VM Client Interface,” on page 15](#)
- ♦ [Section 2.2, “Navigating the Welcome Page,” on page 18](#)
- ♦ [Section 2.3, “Navigating the Inventory Views,” on page 23](#)
- ♦ [Section 2.4, “Saving Changes in the VM Client,” on page 25](#)
- ♦ [Section 2.5, “Understanding the VM Client View and Details Editors,” on page 25](#)

## 2.1 Starting the VM Client Interface

You can run the VM Client on both Windows and Linux workstations:

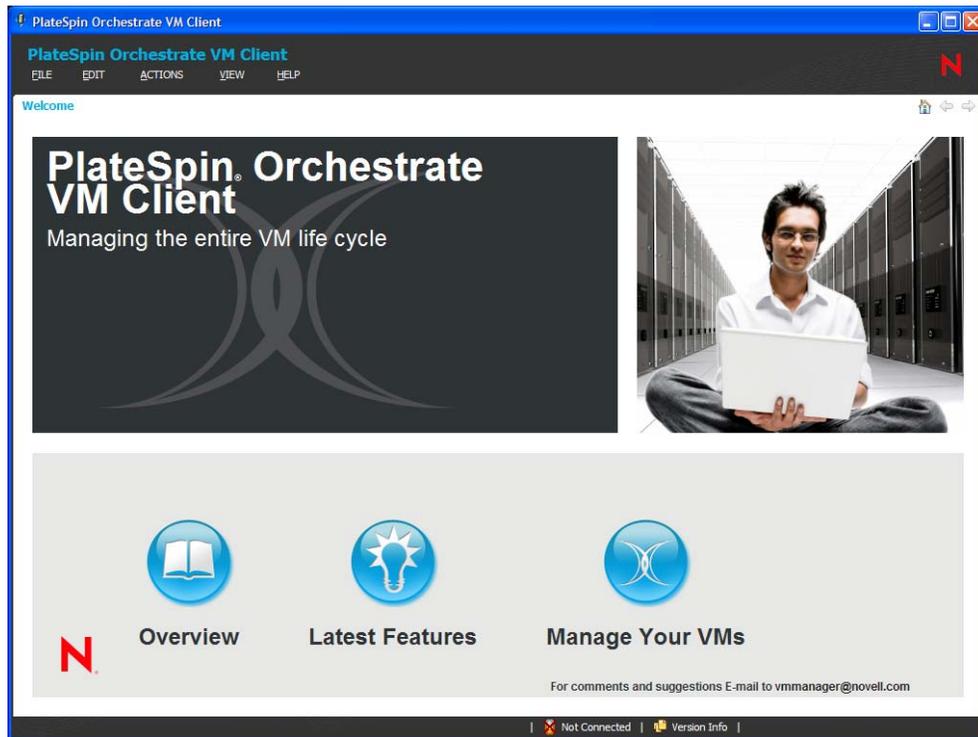
- ♦ [Section 2.1.1, “Starting the Interface on a Windows Workstation,” on page 15](#)
- ♦ [Section 2.1.2, “Starting the Interface on a Linux Workstation,” on page 17](#)

### 2.1.1 Starting the Interface on a Windows Workstation

- 1 On your Windows workstation, double-click the PlateSpin Orchestrate VM Client icon () on your desktop to open the interface.

If you do not have the icon, the executable’s location is determined by where you installed the VM Client software. The default is `C:\Program Files\PlateSpin Orchestrate VM Client 2.0.0\bin\run_nvmm.bat`.

The first time you log in, the following Welcome page is displayed:

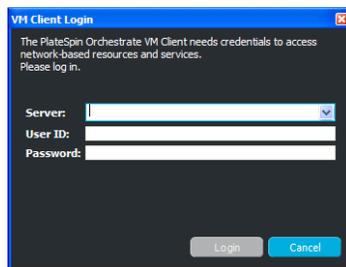


The Welcome page does not display the next time you log in to the VM Client.

The Welcome page can be accessed at any time by clicking *Help > Welcome*.

For more information, see [Section 2.2, "Navigating the Welcome Page,"](#) on page 18.

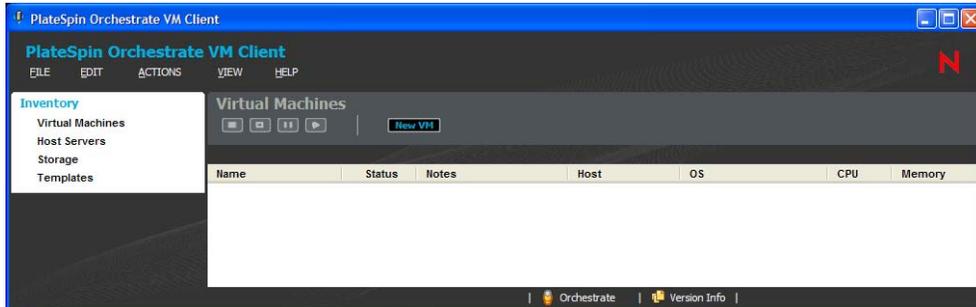
- 2 To close the Welcome page, click *Manage Your VMs* to display the following login dialog box:



- 3 Specify the IP address or DNS name of a PlateSpin Orchestrated Server.

The PlateSpin Orchestrate Server is associated with a virtualization grid, which can be associated with other Orchestrate Servers. Therefore, logging in to one of the Orchestrate Servers allows you to manage all VMs in the grid.

- 4 Enter a username and password that will authenticate you as the administrator of the Orchestrate Server, then click *Login* to view the VM Client interface:



If there is a software compatibility issue that prevents you from logging in, such as upgrading the older software, you must resolve it before you can log in.

## 2.1.2 Starting the Interface on a Linux Workstation

- 1 On your Linux workstation, open a command terminal, change to the `/opt/novell/zenworks/vmmanagement/bin/` directory, then enter the following command:

```
./psvmclient
```

The first time you log in, the Welcome page is displayed. This home page provides the following basic information for the VM Client:

- ♦ **Overview:** Links to a basic overview of the VM Client.
- ♦ **Latest Features:** Gives you links to new features and information about the product. *Latest Features* also highlights any improvements since the previous release of the software.
- ♦ **Manage Your VMs:** Launches the VM Client, where you can begin managing the VMs in your data center. You can also launch the interface by closing the Welcome page.

The main VM Client menu options are also available at the top of the Welcome page, but only those that are applicable are enabled.

The Welcome page can be accessed again at any time by clicking *Help > Welcome*.

- 2 To close the Welcome page, click *Manage Your VMs* to display the Login dialog box.
- 3 Specify the IP address or DNS name of a PlateSpin Orchestrate Server.

The PlateSpin Orchestrate Server is associated with a virtualization grid, which can be associated with other Orchestrate Servers. Therefore, logging in to one of the Orchestrate Servers allows you to manage all VMs in the grid.

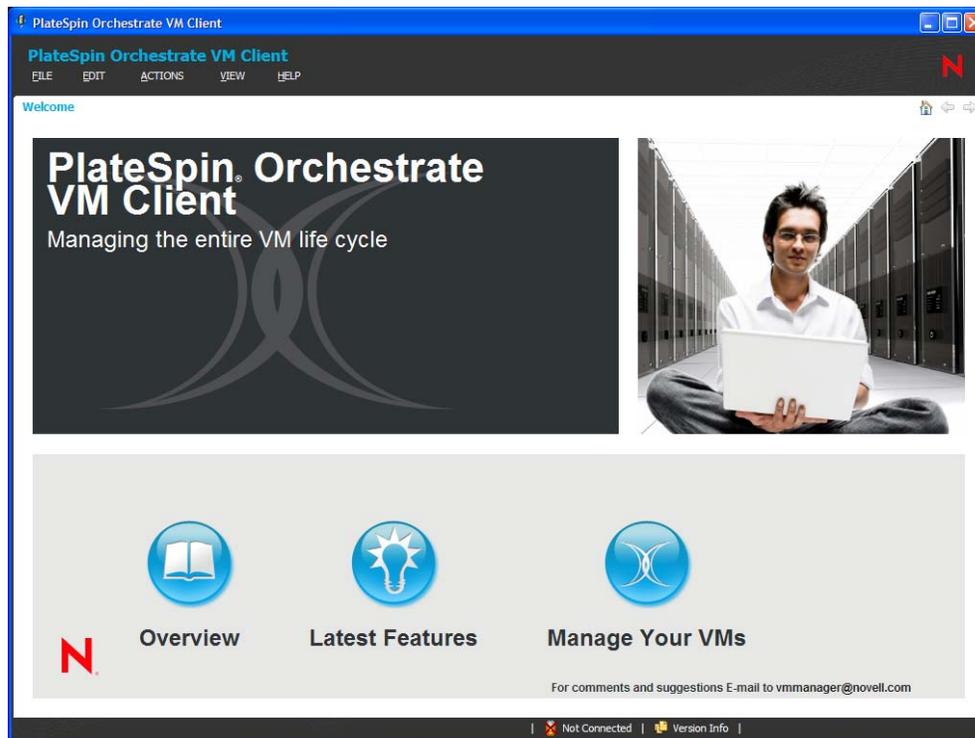
- 4 Enter a username and password that will authenticate you as the administrator of the Orchestrate Server, then click *Login* to view the VM Client interface.

If there is a software compatibility issue that prevents you from logging in, such as upgrading the older software, you must resolve it before you can log in.

## 2.2 Navigating the Welcome Page

The Welcome page provides basic information on the VM Client. By default, it is displayed when you first open the VM Client. Thereafter, you can access the Welcome page using *Help > Welcome*.

**Figure 2-1** PlateSpin Orchestrate the VM Client Welcome Page



The home icon (🏠) in the upper right returns you to this Welcome page. The arrow icons (↶ and ↷) allow you to navigate through previously accessed locations.

Starting the VM Client from your desktop icon or the program's executable does not automatically log you in to the VM Client. This Welcome Page is available without authenticating.

To authenticate, click the *Manage Your VMs* button on the Welcome page to open the VM Client Login dialog box.

The Eclipse-based VM Client GUI allows you to detach any section from the main VM Client window. For example, you can click within the Welcome page and drag it to another location on your desktop. This can also be done with various other sections in the VM Client, such as the **Progress** and **Error Log** views.

The VM Client window is resizable, as are any sections that can be dragged out of the VM Client window. Sections can also be resized within the full VM Client window.

Review the following sections for other navigation tips:

- ◆ [Section 2.2.1, “Welcome Page Main Menu Options,” on page 19](#)
- ◆ [Section 2.2.2, “Welcome Page Buttons,” on page 21](#)

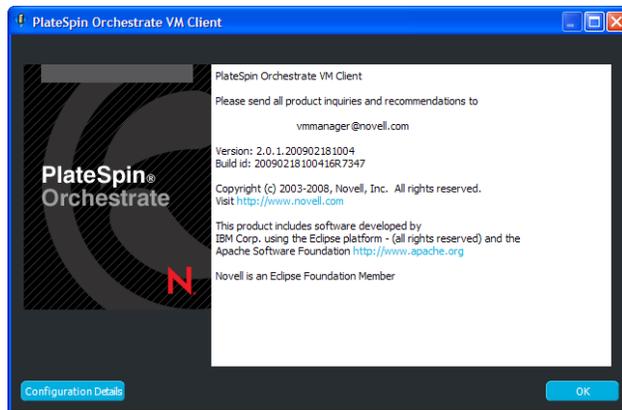
## 2.2.1 Welcome Page Main Menu Options

The main VM Client menu options are available at the top of the Welcome page, but only those that are applicable are enabled:

- ◆ *File > Exit*: Closes the VM Client interface.

Closing the VM Client interface does not affect the status of any VMs, meaning VMs that have processes running, such as cloning, starting, moving, and so on, continue to run those processes. These processes are jobs that are being run on the server hosting the VMs.

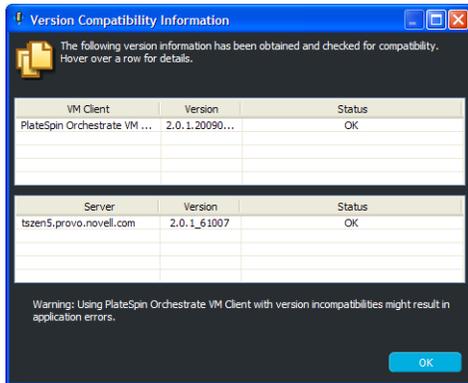
- ◆ *Help > About PlateSpin Orchestrate VM Client*: Displays the following dialog box:



- ◆ The *Configuration Details* button opens the following dialog box where you can copy the text-based configuration information to your operating system's clipboard for further use.

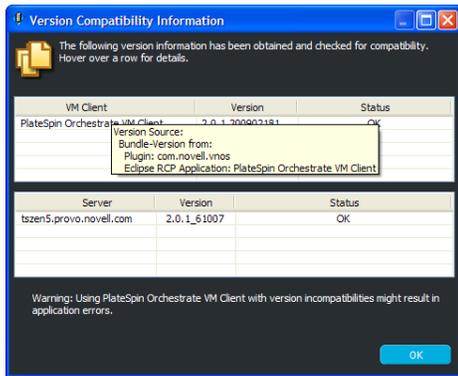


- ◆ *Help > Version Information*: Opens the following dialog box:



This provides software compatibility information that you might need to know when the software is not working correctly.

To view more details for an entry, mouse over its row:



## 2.2.2 Welcome Page Buttons

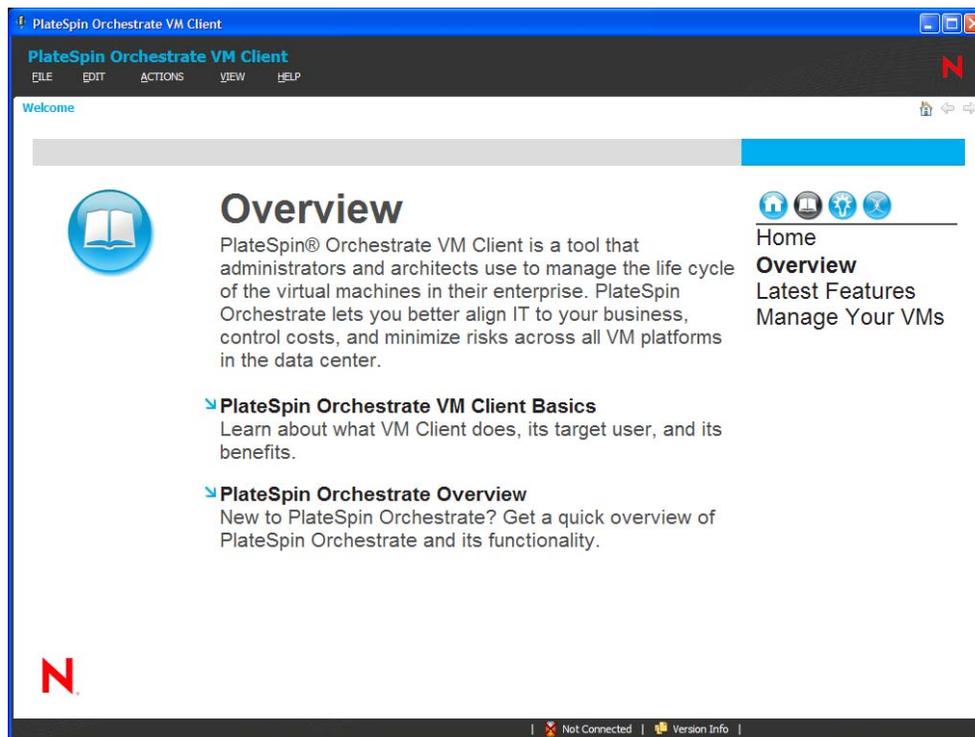
The following buttons are available on the Welcome page:

- ◆ [“Overview” on page 21](#)
- ◆ [“Latest Features” on page 22](#)
- ◆ [“Manage Your VMs” on page 22](#)

### Overview

The *Overview* button provides the following information:

**Figure 2-2** Welcome Page Overview

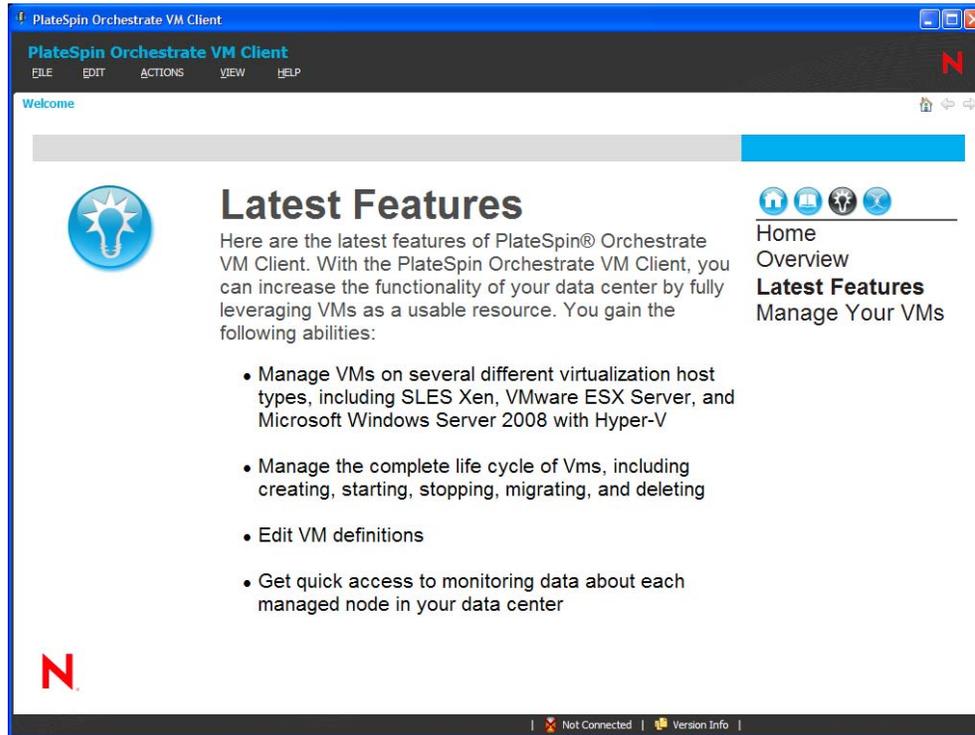


To navigate to the other locations on the Welcome page, you can use the links in the upper right of the Overview page.

## Latest Features

The *Latest Features* button provides the following information:

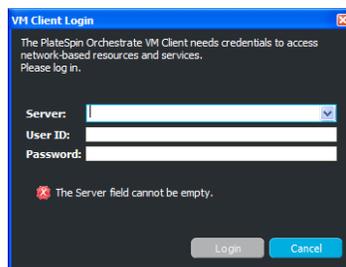
**Figure 2-3** *Welcome Page Latest Features*



## Manage Your VMs

The *Manage Your VMs* button opens the VM Client Login dialog box:

**Figure 2-4** *The VM Client Login Dialog Box*



## 2.3 Navigating the Inventory Views

Review the following to understand the two main VM Client views:

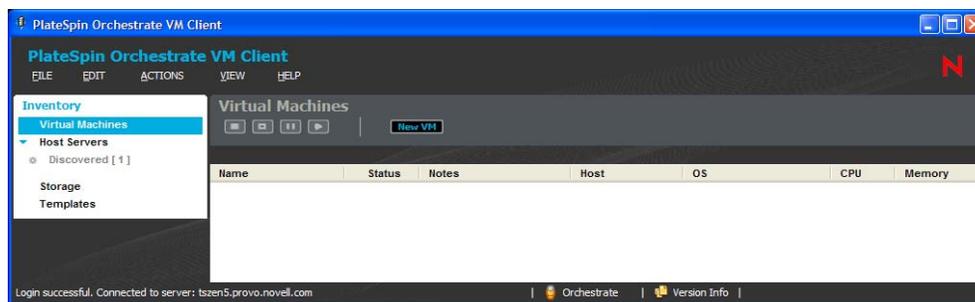
- ◆ [Section 2.3.1, “Inventory Lists,”](#) on page 23
- ◆ [Section 2.3.2, “Inventory Details,”](#) on page 24

For information on the Progress and Error Log views, see [Section 5.17, “Using the Progress View,”](#) on page 73 and [Section 5.16, “Using the Error Log View,”](#) on page 73.

### 2.3.1 Inventory Lists

When you first **log in to the VM Client**, the following view is displayed:

**Figure 2-5** The VM Client Inventory View - Initial View



The Inventory panel on the left contains four links in its navigation area:

- ◆ *Virtual Machines*
- ◆ *Host Servers*
- ◆ *Storage*
- ◆ *Templates*

Each of these four links lists the available items, such as known VMs, known host servers, known repositories, and any templates that you have created from the VMs. By default, the *Virtual Machines* listing is initially displayed.

The following describes the various features for the Inventory view:

- ◆ **Groups:** Each of the four links can have groups under them. The  icon indicates that one or more groups exist. Click the icon to display the groups. For more information on using groups, see [Section 4.8, “Creating and Populating Groups,”](#) on page 59.

When you first log in to the VM Client, if there are newly discovered VMs or host servers that you have not yet accepted into the virtualization grid, the *Virtual Machines* or *Host Servers* sections are automatically expanded to display a *Discovered* group, which you can click to show the candidates in the right panel. For information on managing them, see [Section 3.1, “Registering VM Hosts,”](#) on page 27.

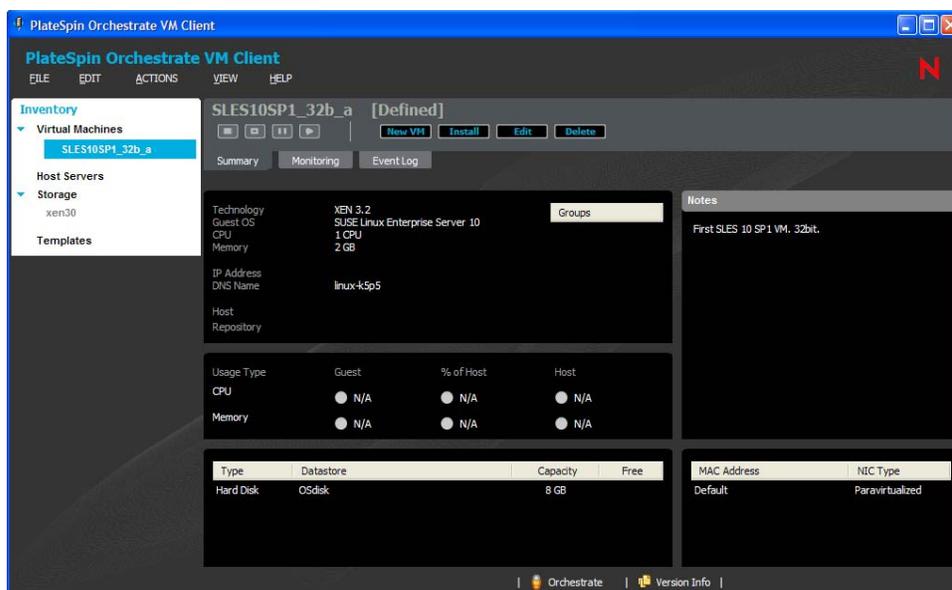
- ◆ **Available Options:** For the items listed in in the right panel for each of the Inventory views, various main menu options, right-click menu options, and function buttons are displayed, not displayed, enabled, or disabled according to the item’s status.

- ◆ **Column Sorting:** You can sort the various columns of information. Simply double-click a column heading to sort it in ascending or descending order. The ▲ or ▼ arrows indicate the column that was last sorted and whether it was sorted in ascending or descending order. For more information, see “[Sorting:](#)” on page 26.
- ◆ **Actions in Progress:** The bottom of the VM Client window displays information applicable to the action in process at the time. This information is also available for the [details views](#).
- ◆ **Logging In/Out:** Mouse over the *Orchestrate* icon to display which host server you are logged in to. You can double-click the icon to log out of that host server and the VM Client. This option is also available for the [details views](#).
- ◆ **Version Information:** Double-click the *Version Info* icon to access the same dialog box that clicking *Help > Version Information* accesses. This option is also available for the [details views](#).

## 2.3.2 Inventory Details

In each of the four Inventory views you can display detailed information by either double-clicking an item or right-clicking the item and selecting *Show Details*. For example:

**Figure 2-6** *Virtual Machines Details Tabs*



The section in the Inventory panel related to the selected item is expanded to show the items listing that was previously displayed to the right, along with tabs containing details for the items.

You can navigate the VMs in the expanded list in the Inventory panel on the left to see the details.

You can navigate through the details for each Inventory view by using the tabs, such as *Summary*, *Monitoring*, and *Event Log* illustrated in [Figure 2-6](#). Each Inventory type has a *Summary* tab plus other tabs particular to its own details. Tab viewing is sticky, meaning the last tab that you viewed for an item is displayed again when you return to that item. For information on using and editing information in each of the details tabs, see the applicable sections in [Section 5.19, “Viewing and Editing VM Client Details,”](#) on page 74.

The *Event Log* tab (for *Virtual Machines*, *Host Servers*, and *Templates*) provides access to details on the status of any jobs that have been run or are running, no matter which administrator initiated the action. For more information, see [“Using Feedback to Manage VMs” on page 65](#).

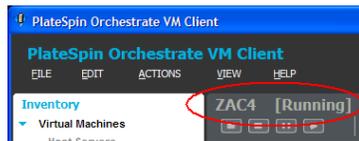
## 2.4 Saving Changes in the VM Client

When you make a change to a field in the VM Client, then exit the field by pressing Enter or changing focus to another field or option, the change made to the field is automatically saved. This is different from the Development Client, where you must click *File > Save* to keep changes to the fields.

## 2.5 Understanding the VM Client View and Details Editors

You can do several things to customize your view of the VM Client. Any changes you make are remembered the next time you start the VM Client.

- ♦ **Tiling:** Multiple editors can be open at the same time. Select the status bar of an editor and drag it to a different position in the VM Client window to create the tiling effect. This can provide quicker access when you revisit an editor.
- ♦ **Maximizing Tab Views:** You can maximize a details tab view by double-clicking the VM or host server’s name that is located just above the buttons. For example:



- ♦ **Dragging:** You can drag views (such as *Progress* or *Error Log*) out of the VM Client window to another position on your desktop. You can also drag views to different locations within the VM Client window.
- ♦ **Returning to Defaults:** To reset any changes that you have made to the positions of views and editors, click *Window > Restore Perspective*, which returns everything to the default organization of the VM Client window.
- ♦ **Keystrokes and Mouse Clicks:** The following functions are available in the Inventory views:

Function	Action
Minimize or maximize of a view or editor in the VM Client window	<ul style="list-style-type: none"> <li>◆ Double-click the status bar</li> <li>◆ Press the Spacebar or Enter key</li> <li>◆ Press the Left-arrow or Right-arrow arrow key</li> </ul>
Go to the next or previous node	<ul style="list-style-type: none"> <li>◆ Press the Left-arrow or Right-arrow arrow key</li> </ul>
Go to the next or previous sibling node	<ul style="list-style-type: none"> <li>◆ Press the Tab or Shift+Tab keys</li> </ul>
Open further details on a clickable details item	<ul style="list-style-type: none"> <li>◆ Double-click a details item, such as a line in the <i>Event Log</i> tab that has the  icon, to view further details on the event entry.</li> </ul>

- ◆ **Sorting:** Most lists default to sorting alphabetically by the first column. You can sort any column in an ascending or descending order by simply double-clicking the column. The column last sorted displays a sorting arrow (▲ or ▼).

Sorting in the VM Client provides the following benefits:

- ◆ When new items are added, they are automatically sorted into the listing.
- ◆ Sorting on the *Status* column for VMs groups them by their status type in the following order:



- ◆ Sorting on the *Host* column for VMs is a quick way to determine per host server which VMs are currently running or in a progress state.
- ◆ Sorting on the *CPU* or *Memory* columns for VMs is a quick way to determine your host servers' utilization if the Orchestrate and Monitoring Agents are installed on the host servers.

# Creating and Setting Up Virtual Machines

# 3

Use the information in the following sections to set up the PlateSpin® Orchestrate VM Client and to create virtual machines (VMs):

- ◆ Section 3.1, “Registering VM Hosts,” on page 27
- ◆ Section 3.2, “Discovering Registered VM Hosts,” on page 29
- ◆ Section 3.3, “Discovering VMs of the Registered VM Hosts,” on page 29
- ◆ Section 3.4, “Registering VMs,” on page 30
- ◆ Section 3.5, “Creating a Xen VM,” on page 31
- ◆ Section 3.6, “Installing a VM,” on page 38
- ◆ Section 3.7, “Installing the PlateSpin Orchestrate Agent on a VM,” on page 38
- ◆ Section 3.8, “Installing the Monitoring Agent on a VM,” on page 39
- ◆ Section 3.9, “Creating a Template from a VM,” on page 40
- ◆ Section 3.10, “Adding Repositories,” on page 41
- ◆ Section 3.11, “Associating Repositories,” on page 43
- ◆ Section 3.12, “Removing Repositories,” on page 43

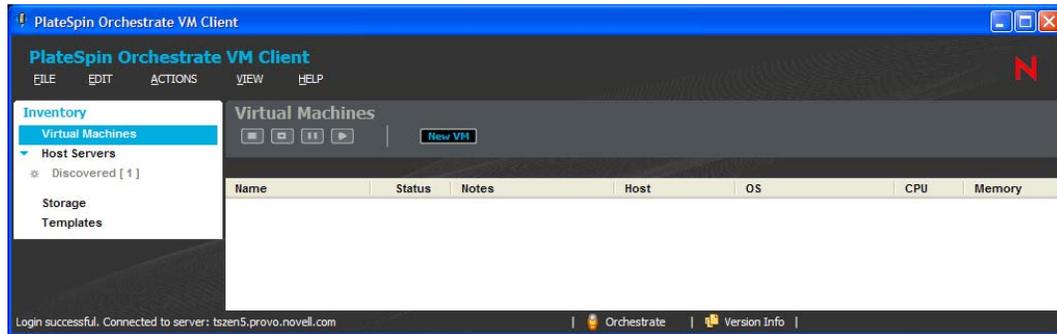
## 3.1 Registering VM Hosts

After being installed on a computing node, having its credentials defined, and associating itself with the computing node, the Orchestrate Agent begins broadcasting the availability of its host as a potential computing resource. A host is defined as a machine running a **supported hypervisor agent**.

However, before the Orchestrate Server can allow an agent to authenticate and establish ongoing communication, you need to create a resource account for the agent on the Orchestrate Server. When this account is created or “registered,” the agent’s host node can be discovered and recognized as a computing resource that can perform the **jobs** assigned to it. In the Orchestrate Development Client, you can choose to automatically or manually register a VM host server to the Orchestrate Server by configuring the *Auto Register Agents* option (the *Authentication* tab of the grid > the Resources panel). If you choose to manually register to the Orchestrate Server, the unregistered VM host servers are displayed in the VM Client.

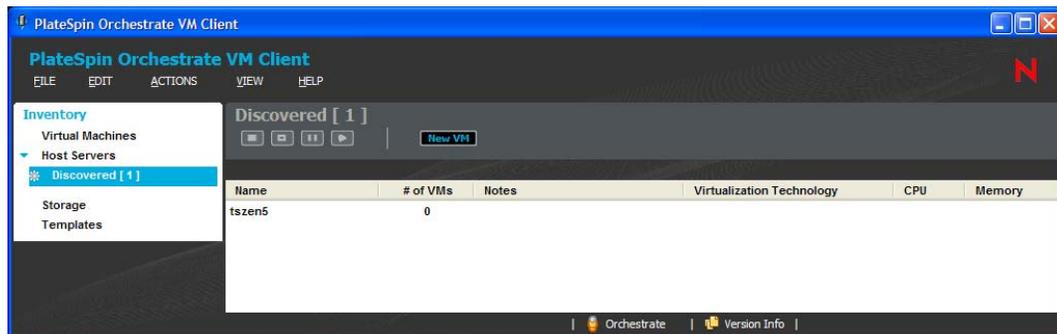
To register the VM host servers:

- 1 In the **VM Client**, double-click the *Host Servers* view.



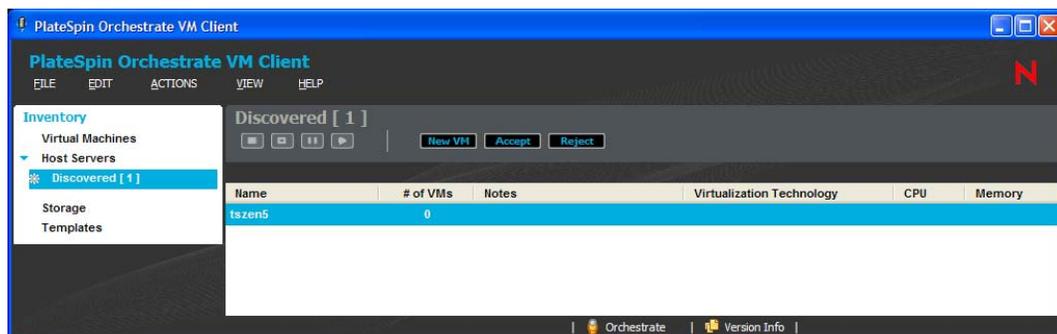
*Discovered (#)* is displayed below *Host Servers*, where # represents the number of host servers that are waiting to be registered with the Orchestrator Server.

- 2 Click *Discovered (#)* to display the unregistered host servers in the details section on the right:



You can accept (see [Step 3](#)) or reject (see [Step 4](#)) any of the listed host servers. The host servers that you accept become part of your virtualization grid.

- 3 Accept an unregistered host server in one of the following ways:
  - ◆ Select the server in the details view, then click the **Accept** button.
  - ◆ Right-click the server in the details view, then select *Accept*.
  - ◆ Select the server in the details view, then click *Actions > Accept*.



You can choose to register multiple host servers at a time.

After the host server has been registered with the Orchestrate Server, it is listed with its operating system.

If the host server doesn't become registered after a few minutes, run the *Discover Hosts* job from the *Action* menu.

4 (Optional) To reject the registration of a host server, do one of the following:

- ♦ Select the server in the details view, then click the **Reject** button.
- ♦ Right-click the server in the details view, then select *Reject*.
- ♦ Select the server in the details view, then click *Actions > Reject*.

The unregistered server is removed from the Host Servers listing.

The next time that an unregistered host server attempts to register to the Orchestrate Server, the rejected servers are again listed as candidates for acceptance into the virtualization grid.

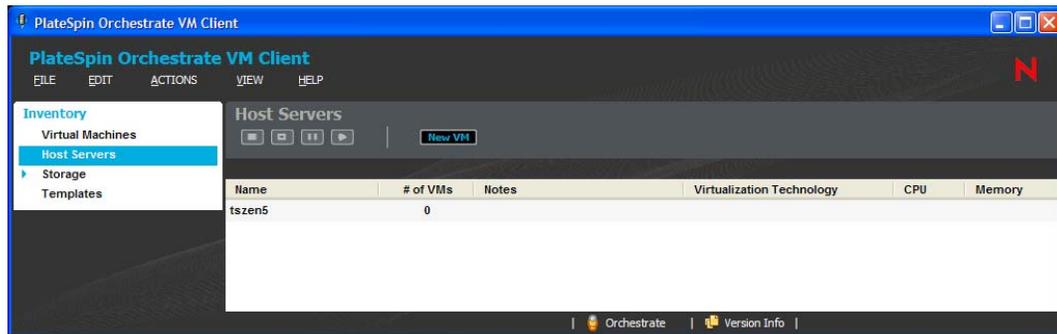
## 3.2 Discovering Registered VM Hosts

When you first install the VM Client, it is not aware of all of the possible registered VM hosts in the virtualization grid that you log in to. You should discover these hosts before continuing to use the product. You should discover hosts before attempting to discover VMs so that any VMs belonging to a particular host can be discovered.

Before discovering the VMware technology-based hosts, ensure that the appropriate policies have been configured in the Orchestrate Development Client. For more information on configuring the policies, see “[Configuring Policies for VM Provisioning Adapters](#)” in the *PlateSpin Orchestrate 2.0 Virtual Machine Management Guide*.

To discover the existing registered VM hosts:

- 1 In the **VM Client**, click *Actions > Discover Hosts*.



All the registered VM hosts that are in the grid and that have hypervisors installed are displayed.

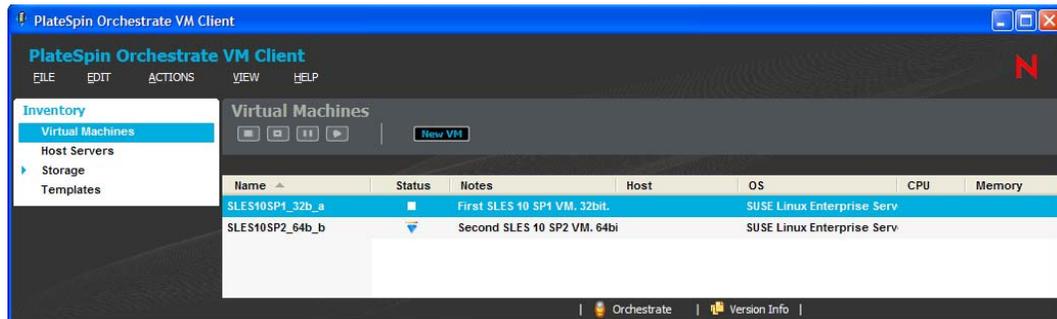
## 3.3 Discovering VMs of the Registered VM Hosts

You can use the VMs you have previously created or the VMs you have built by using other hypervisors in your data center. To make use of these VMs, you need to run a **discover job**. You use the VM Client to detect the VMware ESX, Microsoft Hyper-V, VMware Server, VMware Virtual Center, and SUSE® Xen VM host machines. After you have detected the host machines, you can detect the VMs contained on the host machines and in other VM storage repositories.

Before discovering the VMware technology-based VMs, ensure that the appropriate policies have been configured in the Orchestrate Development Client. For more information on configuring the policies, see “[Configuring Policies for VM Provisioning Adapters](#)” in the *PlateSpin Orchestrate 2.0 Virtual Machine Management Guide*.

To discover the VMs belonging to the registered VM Hosts:

- 1 In the **VM Client**, click *Actions > Discover Virtual Machines*.



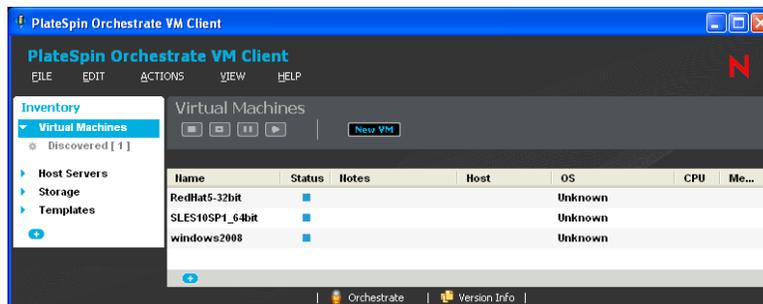
All VMs of the registered host servers in the virtualization grid are displayed and automatically registered.

### 3.4 Registering VMs

In the Orchestrate Development Client, you can choose to automatically or manually register a VM to the Orchestrate Server by configuring the *Auto Register Agents* option (the *Authentication* tab of the grid > the Resources panel). If you choose to manually register to the Orchestrate Server, the unregistered VMs are displayed in the VM Client.

To manually register the VMs:

- 1 In the **VM Client**, double-click the *Virtual Machines* view.



*Discovered (#)* is displayed below *Virtual Machines*, where # represents the number of VMs that are waiting to be registered with the Orchestrate Server.

- 2 Click *Discovered (#)* to display the unregistered VMs in the details section on the right. You can accept (see [Step 3](#)) or reject (see [Step 4](#)) any of the listed VMs that were discovered. The VMs that you accept become part of your virtualization grid.
- 3 Accept a discovered VM in one of the following ways:
  - ♦ Select the VM in the details view, then click the **Accept** button.

- ♦ Right-click the VM in the details view, then select *Accept*.
- ♦ Select the VM in the details view, then click *Actions > Accept*.

You can choose to register multiple VMs at a time.

After a VM has been registered, it is listed with its operating system.

Because it can take a few minutes for this acceptance process to be completed, do not proceed with using the VMs until it has.

**4** (Optional) To reject a VM, do one of the following:

- ♦ Select the VM in the details view, then click the **Reject** button.
- ♦ Right-click the VM in the details view, then select *Reject*.
- ♦ Select the VM in the details view, then click *Actions > Reject*.

The unregistered VM is removed from the Virtual Machines listing.

The next time that an unregistered VM attempts to register to the Orchestrate Server, the rejected VMs are again listed as candidates for acceptance into the grid.

## 3.5 Creating a Xen VM

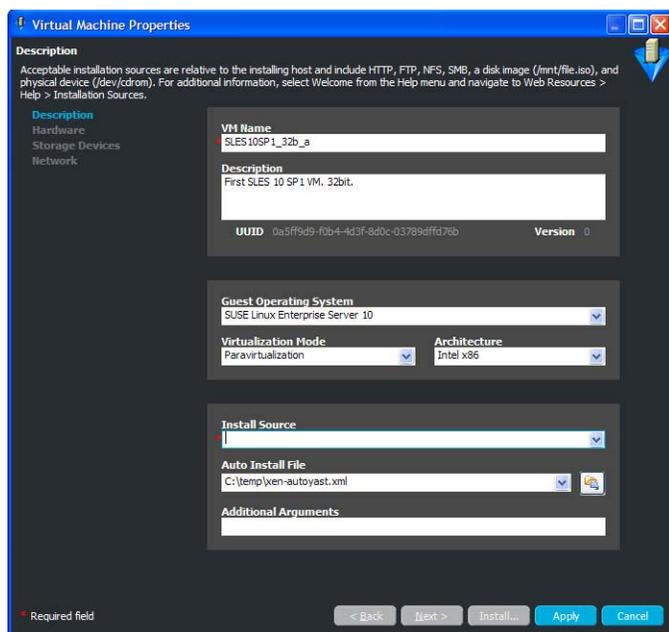
Refer to the appropriate third-party documentation for how to create VMs for VMware ESX Server and Microsoft Windows Server 2008 with Hyper-V. To bring them into the virtualization grid, see [Section 3.3, “Discovering VMs of the Registered VM Hosts,” on page 29](#).

To create a new Xen VM for the virtualization grid that you are logged in to, run the Virtual Machine Properties Wizard:

**1** In the **VM Client**, do one of the following to start the wizard:

- ♦ Click the **New VM** button in any view.
- ♦ Click *File > New VM* in any view.
- ♦ Press Ctrl+N in any view.
- ♦ Right-click any listed item in any view, then select *New VM*.

The following wizard page is displayed:



## 2 Fill in the fields to define the basic information for the VM:

- ◆ **VM Name:** (Required) Name for the VM.

This name applies to all versions of the VM and must be unique. It cannot already exist in either the *Virtual Machines* or *Templates* listings.

The name cannot be longer than 100 characters and cannot start with the letters “xen.”

- ◆ **Description:** Note the purpose and any other specifics you might require of the VM. These notes apply to all versions of this VM.
- ◆ **UUID:** Displays the UUID for the VM. This field is view-only.
- ◆ **Guest Operating System:** The operating system you want to install on your VM.
- ◆ **Virtualization Mode:** *Full Virtualization* means that all the hardware components are emulated by the hypervisor. *Paravirtualization* means that the operating system directly accesses elements of the hardware, including processor, display, memory, network interface card, and hard drive, and uses the hypervisor to emulate other hardware interactions. Paravirtualization performs much faster and with better hardware efficiency than full virtualization.

---

**NOTE:** Refer to the SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux\* (RHEL) documentation to verify supported platforms in each virtualization mode.

---

- ◆ **Architecture:** The type of processor you want to have as the virtual processor for the VM.
- ◆ **Install Source:** (Required) The path to the installation media for your guest operating system. For a fully virtualized VM, the install source must be an ISO. For a paravirtualized VM, an ISO or network install source is required. For more information, see [Appendix A, “VM Installation Sources,” on page 93](#).
- ◆ **Auto Install File:** This field applies only to paravirtual installations. The file listed here is used to populate the answers needed during the installation of the guest operating system. Some examples include time zone, language, firewall, and other settings given in response

to the user interactive areas of the installation. An AutoYaST file listed here can contain network settings that are applied during the VM Builder job. For RHEL, the file is a kickstart file.

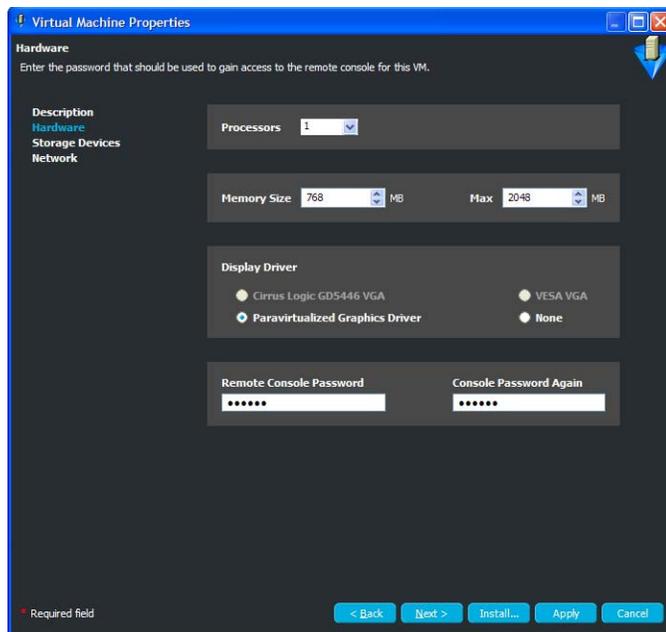
---

**IMPORTANT:** The auto install file must reside locally on the machine where you are running the VM Client. It cannot be used from a host machine or another machine in your network.

---

- ♦ **Additional Arguments:** This field applies only to paravirtual installations. These are Linux install arguments, which are operating system parameters similar to the information entered in a GRUB boot loader when running on bare hardware in SLES and SUSE Linux Enterprise Desktop (SLED). Additional network installation information can be added here, which is used during the build process to find the installation source for the VM build process.

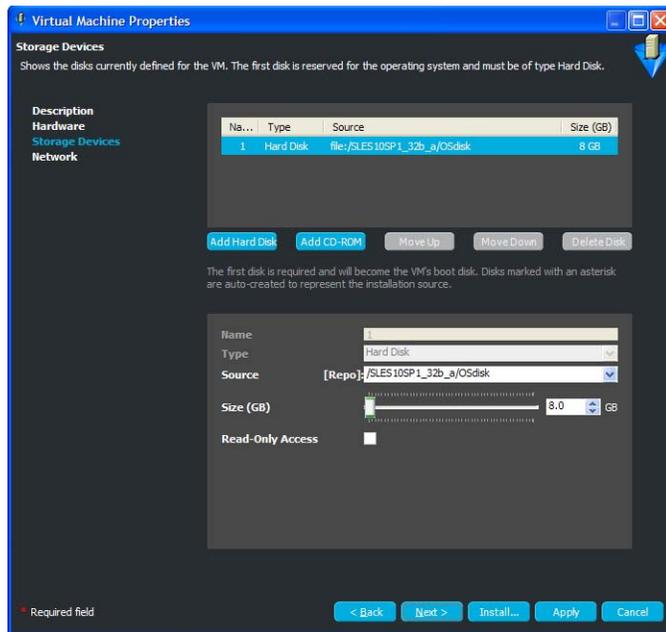
3 Click *Next* to display the following wizard page:



4 Fill in the fields to define the kind of hardware to be virtualized for the VM:

- ♦ **Processors:** The number of virtual processors the VM has.
- ♦ **Memory Size:** The amount of virtual RAM the VM has, or paravirtualized memory that can be used by the VM if it is paravirtualized.
- ♦ **Max:** The amount of virtual RAM the VM has, or paravirtualized memory that can be used by the VM if it is paravirtualized.
- ♦ **Display Driver:** The properties of the VM's display. If you have chosen paravirtualization for this VM, your only option is *Paravirtualized*. If you have chosen full virtualization, you can select between *VESA* or the on-board VGA. If you do not need a display, you can turn off the display on a fully virtualized machine.
- ♦ **Remote Console Password:** The password for remotely controlling the VM. Use the two fields to verify that it is typed correctly.

5 Click *Next* to display the following wizard page:



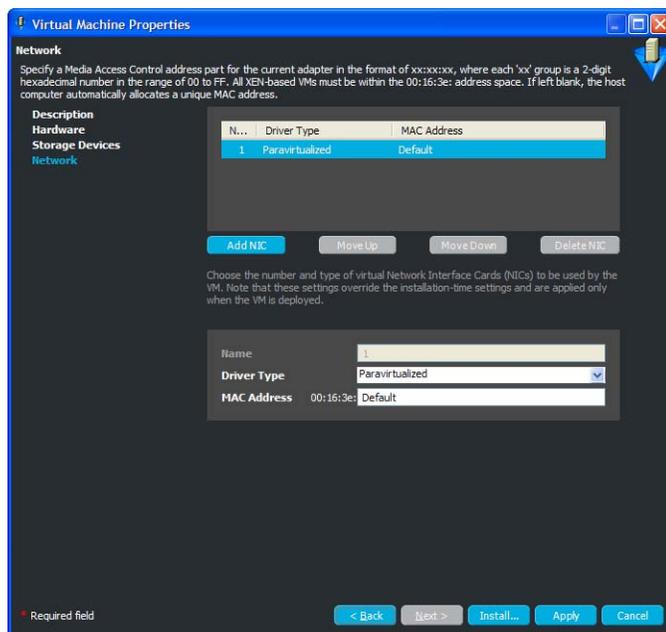
The Storage Devices page is for defining the hard drives and CD or DVD drives that you want to include in the VM profile. The first device must always be a hard disk that cannot be removed unless another hard disk has been put into the first slot. This is the only disk (Disk 0) that is stored and version controlled. All other disks are virtual and are created upon launching the VM. You can add any other storage devices you want to the VM and move them according to priority in the list.

6 To fill in the storage device information, select a task from the following table:

Tasks	Steps	Additional Details
Adding a hard disk to your VM	<ol style="list-style-type: none"> <li>Click <i>Add Hard Disk</i>.</li> <li>Fill in the following fields: <ul style="list-style-type: none"> <li><b>Name:</b> The name is the number in the order the drive appears in its VM image listing.</li> <li><b>Type:</b> Select the type of hard disk you want to create on the VM: <ul style="list-style-type: none"> <li>Hard Disk</li> <li>iSCSI</li> <li>N_Port ID Virtualized Storage</li> <li>Network Attached Storage</li> </ul> </li> <li><b>Source:</b> The repository on the hard disk of the VM.</li> <li><b>Size (GB):</b> The VM's hard drive size, ranging from 0.25 GB (250 MB) to 4 TB. This is the maximum allowable size of the VM.</li> <li><b>Read Only Access:</b> Select this option to give the hard drive Read-Only access. By default, hard drives are writable.</li> <li><b>Fabric ID:</b> (Available only for the N_Port ID virtualized storage). Select or specify a <b>fabric ID</b> that uniquely identifies a SAN repository. The <i>Fabric ID</i> drop-down list is populated with fabric IDs of the discovered SAN repositories.</li> <li><b>World Wide Name:</b> (Available only for the N_Port ID virtualized storage). Select or specify a unique <b>World Wide Name</b>. The World Wide Name value is automatically assigned for the NPIV disks. You can edit it, if required.</li> <li><b>Target ID:</b> (Available only for the N_Port ID virtualized storage). Specify the <b>target ID</b> of the SAN repository.</li> <li><b>LUN:</b> (Available only for the N_Port ID virtualized storage). Specify a <b>LUN</b> value between 0 and 254.</li> </ul> </li> <li>Click <i>Apply</i>.</li> </ol>	<p>The following restrictions apply per VM:</p> <ul style="list-style-type: none"> <li>Limit of 4 fully virtualized disks (hda-hdd).</li> <li>Limit of 16 paravirtualized disks (xvda-xvdp).</li> </ul> <p>While adding an additional NPIV disk to a VM, ensure that the disk values for Fabric ID, World Wide Name, Target ID, and LUN are unique for each NPIV disk of the VM.</p> <p>Creation of a block-based VM is not supported for NPIV.</p>

Tasks	Steps	Additional Details
Adding a CD-ROM or DVD ROM to your VM	<ol style="list-style-type: none"> <li>Click <i>Add CD-ROM</i>.</li> <li>Fill in the following fields: <ul style="list-style-type: none"> <li><b>Name:</b> The name is the number in the order the drive appears in its VM image listing.</li> <li><b>Type:</b> Select the type of CD-ROM you want to create on the VM: <ul style="list-style-type: none"> <li>◆ CD-ROM</li> <li>◆ DVD</li> </ul> </li> <li><b>Source:</b> The repository on the hard disk of the VM.</li> <li><b>Size (GB):</b> The size is zero and cannot be changed.</li> <li><b>Read Only Access:</b> These drives can only have Read-Only access.</li> </ul> </li> <li>Click <i>Apply</i>.</li> </ol>	
Rearranging the order in which each storage device is accessed	<ol style="list-style-type: none"> <li>Select the storage device whose order you want to change.</li> <li>Click <i>Move Up</i> or <i>Move Down</i>.</li> </ol>	The first storage device in the list must be a hard disk, which is the only one actually stored in the VM image.
Deleting a storage device	<ol style="list-style-type: none"> <li>Select the storage device you want to delete.</li> <li>Click <i>Delete</i>.</li> </ol>	Removes the selected storage device from the table.

7 Click *Next* to display the following wizard page:



The network devices are added, deleted, and defined on the Network page, where you add specific information about your network connections or intended network connections for your VM. You can define more detailed NIC settings in your auto-install files or when performing a guest installation.

**8** To fill in the NIC information, select a task from the following table:

Tasks	Steps
Adding a NIC to your VM	<ol style="list-style-type: none"> <li>1. Click <i>Add NIC</i>.</li> <li>2. Fill in the following fields: <ul style="list-style-type: none"> <li>◆ <b>Name:</b> Name of the NIC card. This name is used to help you to recognize each NIC.</li> <li>◆ <b>Driver Type:</b> Indicates whether the VM has direct access to the NIC. If the NIC should only interface with the hypervisor, you can select the specific type of full virtualization. You should have a paravirtualized NIC with a paravirtualized VM.</li> <li>◆ <b>MAC Address:</b> Unless you have a specific MAC address you want to assign to this VM, leave this set to <i>Default</i>, and Xen assigns it a MAC address when you launch the VM.</li> </ul> </li> <li>3. Click <i>Apply</i>.</li> </ol>
Rearranging the order in which the NICs are accessed	<ol style="list-style-type: none"> <li>1. Select the storage device whose order you want to change.</li> <li>2. Click <i>Move Up</i> or <i>Move Down</i>.</li> </ol>
Deleting a NIC from your VM	<ol style="list-style-type: none"> <li>1. Select the NIC you want to remove.</li> <li>2. Click <i>Delete NIC</i>.</li> </ol>

**9** To complete the wizard, click one of the following buttons:

- ◆ **Install:** Lists the VM in the *Virtual Machine* view, then immediately installs the configured VM.

A VNC login dialog box is displayed so that you can view the progress of the installation from the VM's console. There is a 30-second timeout while the VNC sets up the parameters of the VM.

You can cancel the use of VNC; however, if user interaction is necessary, you need to use VNC to help complete the installation. You can start the VNC for a VM being installed by right-clicking the VM and clicking *Show Console*.

- ◆ **Apply:** Applies the wizard configuration and saves the new VM in the *Virtual Machines* view, but does not install the VM.

Focus is returned to the *Virtual Machines* view after you exit the wizard, with the new VM selected in the list.

To install the VM at a later time, see [Section 3.6, "Installing a VM," on page 38](#).

You can click *Back* to make changes before installing or applying the VM definition.

## 3.6 Installing a VM

To install an existing VM definition:

- 1 (Optional) In the **VM Client**, click the *Virtual Machine* view, right-click the VM that you want to install, then select *Edit*.

If you have any changes for the VM, complete those by using the *Next* and *Back* buttons.

For more information on editing the wizard pages, see [Section 4.1, “Editing VMs or Templates,” on page 45](#).

- 2 To install the VM, do one of the following:
  - ♦ While you are in the Edit VM wizard, click *Install*.
  - ♦ In the *Virtual Machines* view, select the defined VM (it shows  in the *Status* column), then click the  button.
  - ♦ In the *Virtual Machines* view, right-click the defined VM (it shows  in the *Status* column), then select *Install*.
- 3 Proceed with one of the following:
  - ♦ If the installation is proceeding from an installation source that requires user input, the VNC is automatically started. Log in with the VNC password to view and interact with the installation as it progresses on the VM’s console, then continue with [Step 4](#).
  - ♦ If the installation is using YaST, the VM is automatically installed without user input. You can either manually use VNC (click *View > Show Console*) to view the installation’s progress, or you can continue with [Step 4](#).
- 4 You can view the installation progress in any of the following ways:
  - ♦ Double-click the VM being installed, click the *Event Log* tab, then double-click the entry related to the installation process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
  - ♦ Observe messages at the bottom of the VM Client interface.
  - ♦ Click *View > Show Progress View* to open the Progress window.
  - ♦ View its progress in the *Jobs* tab of the **Development Client**.
- 5 If you want to cancel the installation before it finished, do one of the following:
  - ♦ Select the VM being installed, then click *Edit > Cancel Install*.
  - ♦ Right-click the VM being installed, then select *Cancel Install*.

The VM’s status is returned to being configured but not installed.

## 3.7 Installing the PlateSpin Orchestrate Agent on a VM

For instructions on installing the agent on a host server, see “[Installing the Orchestrate Agent Only](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

- 1 In the **VM Client**, click the *Virtual Machines* view, right-click the VM where you want to install the agent, then select *Shut Down*.

The VM cannot be running during installation of the agent because the VM’s image is modified.

The agent cannot be installed on a template.

- 2 After the VM has completed shutting down, select the VM again, then do one of the following:
  - ♦ Click *Actions > Install Agent*.
  - ♦ Right-click the VM, then select *Install Agent*.

The agent installation begins.

- 3 You can view the installation progress in any of the following ways:
  - ♦ Double-click the VM being installed, click the *Event Log* tab, then double-click the entry related to the installation process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
  - ♦ Observe messages at the bottom of the VM Client interface.
  - ♦ Click *View > Show Progress View* to open the Progress window.
  - ♦ View its progress in the *Jobs* tab of the **Development Client**.

You can tell which VMs have the Orchestrate Agent installed by whether pie charts are being displayed in the *CPU* and *Memory* columns in the *Virtual Machines* view, but only if the VM is running.

## 3.8 Installing the Monitoring Agent on a VM

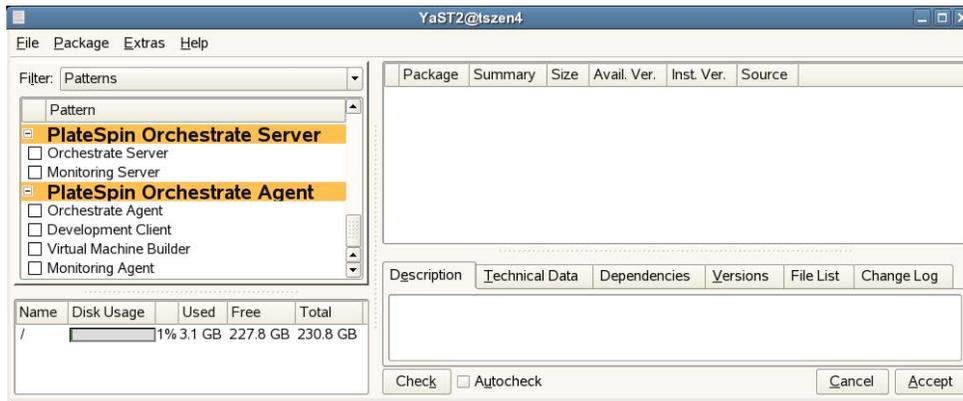
In order to use the **Monitoring feature** on VMs, you must manually install the Monitoring Agent on them.

The Monitoring Agent is usually installed on host servers during installation of the PlateSpin Orchestrate Server product. If it is not installed, you can install it manually by using the following instructions.

The Monitoring Agent does not work on Windows devices, so the following instructions are only for Linux VMs.

- 1 Download the appropriate PlateSpin Orchestrate Server ISO (32-bit or 64-bit) to an accessible network location.
- 2 (Optional) Create a DVD ISO (32-bit or 64-bit) that you can take with you to the machine where you want to install it.
- 3 Install the Monitoring Agent software:
  - 3a Log in to the target VM or host server as `root`, then open YaST2.
  - 3b In the YaST Control Center, click *Software*, then click *Add-on Product* to display the Add-on Product Media dialog box.
  - 3c In the Add-on Product Media dialog box, do one of the following to select the ISO media for installation:
    - ♦ Select *DVD*, click *Next*, insert the DVD, then click *Continue*.
    - ♦ Select *Local Directory*, click *Next*, select the ISO Image check box, browse to and select the ISO in the file system, then click *OK*.

**3d** Read and accept the license agreement, then click *Next* to display the following page:



**3e** In YaST2, click the *Filter* drop-down menu, then select *Patterns* to display the install patterns available on the PlateSpin Orchestrate ISO.

**3f** Select only the *Monitoring Agent* installation pattern.

This pattern installs the Ganglia Agent on each monitored node, which collects performance metrics and sends the data to the Orchestrate Monitoring Server. This pattern includes packages that help you configure the Agent.

**3g** Click *Accept* to install the package.

## 3.9 Creating a Template from a VM

Templates allow you to quickly clone new VMs that are based on the template's configuration. You create a template from an existing VM, but it is created with some default settings, so it is not identical to the VM.

You can modify a template's configuration, if necessary. However, after you have cloned VMs from a template, the template configuration can no longer be modified without first deleting the VMs that were cloned from it. Therefore, when you create the template, plan to make all of your modifications before using it to clone new VMs.

To create a template:

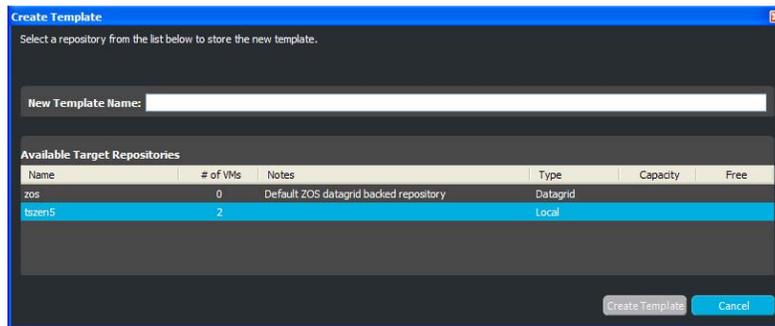
**1** In the **VM Client**, click the *Virtual Machines* view, right-click the VM that you want to use to create the template, then select *Shut Down*.

The VM cannot be running during creation of the template because the VM's image is modified.

**2** In the **VM Client**, click the *Virtual Machine* view, select the VM that you want to use to create the template, then do one of the following:

- ◆ Click the **Template** button.
- ◆ Click *File > Create Template*.
- ◆ Right-click the VM, then select *Create Template*.

The following dialog box is displayed:



3 Specify a template name in the *New VM Name* field.

4 Select the repository for the template.

Template files can be stored anywhere that provides access to them. The known repositories are displayed.

5 Select the host for the template.

The VM hosts connected to the selected repository are displayed. Templates must be associated with a VM host.

6 Click *Create Template*.

You can view the progress in any of the following ways:

- ◆ Double-click the VM from which the template is being created, click the *Event Log* tab, then double-click the entry related to the creation process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
- ◆ Observe messages at the bottom of the VM Client interface.
- ◆ Click *View > Show Progress View* to open the Progress window.
- ◆ View its progress in the *Jobs* tab of the **Development Client**.

The new template is listed in the *Templates* view.

7 To modify the template to your desired specifications, continue with [Section 4.1, “Editing VMs or Templates,” on page 45](#).

After editing the template, you can then use it to clone new VMs. For more information, see [Section 4.6, “Cloning Templates,” on page 57](#).

## 3.10 Adding Repositories

You can add new repositories for a VM’s files, including its image. *Local* is the default, but the path information is blank and must be filled in with known data.

1 On the server where you plan to specify a new repository, create the path that you plan to use and make a note of it.

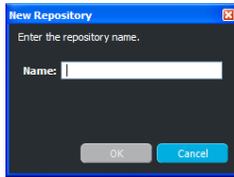
When providing this path information in [Step 6](#), the fields do not provide browsing.

2 In the **VM Client**, click the *Storage* view, then do one of the following:

- ◆ Click the  button.

- ◆ Click *File > New Repository*.
- ◆ Press Ctrl+R.

The following dialog box is displayed:

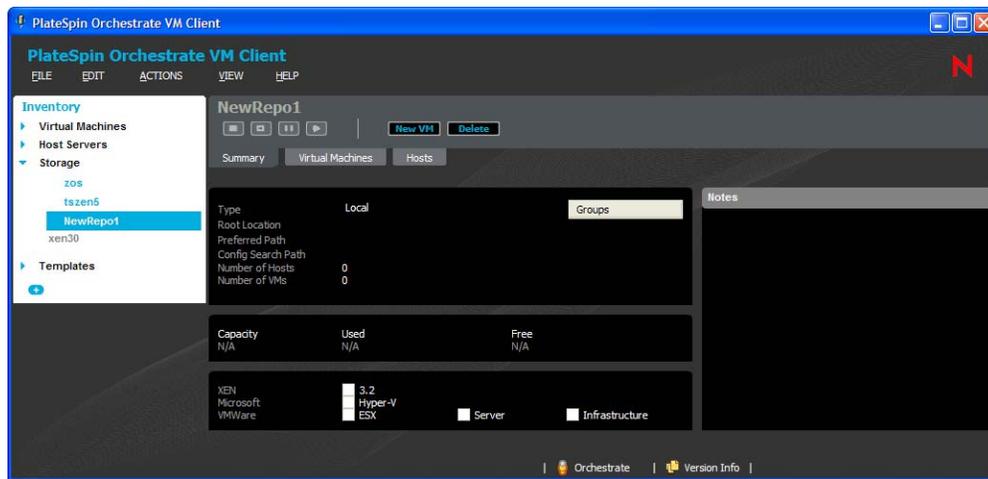


- 3 Specify a name for the repository, then click *OK*.

A default local repository is created, using default paths and other information.

- 4 Right-click the new repository, then select *Show Details*.

The following is displayed:



- 5 On the *Summary* tab, from the drop-down list for the *Type* field, select one of the following:
  - ◆ Network Attached Storage
  - ◆ Local

The following also display, but are not selectable:

- ◆ Storage Area Network
- ◆ Datagrid
- ◆ Virtual

- 6 Fill in the fields:

**Root Location:** For example, “/” (root) for *Local*.

**Preferred Path:** For example, `var/lib/xen/images` for *Local*. This path does not start with a forward slash. The root of the path is provided in the *Root Location* field.

**Config Storage Path:** For example, `/etc/xen/vm` for *Local*. Be sure to include the root of the path, such as the forward slash.

The *Number of Hosts* and *Number of VMs* fields display only hosts and VMs that are currently using this repository.

- 7 (Optional) To enable any of the following for this repository, click the check box. Unavailable options are dimmed.
- ◆ XEN 3.2
  - ◆ Microsoft Hyper-V
  - ◆ VMWare ESX
  - ◆ Server
  - ◆ Infrastructure
- 8 (Optional) Provide any useful information in the *Notes* field.  
When you leave the field, the information is saved.

## 3.11 Associating Repositories

For information on creating new repositories, see [Section 3.10, “Adding Repositories,”](#) on page 41.

- 1 In the **VM Client**, click the *Host Servers* view, right-click a host machine, then select *Show Details*.
- 2 Right-click anywhere in the storage section, then select *Associate Repository*.

The following dialog box is displayed:



- 3 Select a repository, then click *OK*.

The repository is listed with the others for the selected host server. The VMs using this repository can now be run from this host server.

## 3.12 Removing Repositories

You can remove repositories from being associated with a host server, or you can remove a repository altogether:

- ◆ [Section 3.12.1, “Removing a Repository from a Host Server,”](#) on page 44
- ◆ [Section 3.12.2, “Removing a Repository from the VM Client,”](#) on page 44

### 3.12.1 Removing a Repository from a Host Server

---

**IMPORTANT:** When you remove the repository from association with the host server, the VMs stored in the repository can no longer be managed through this host server.

---

- 1 In the **VM Client**, click the *Host Servers* view, right-click a host server, then select *Show Details*.
- 2 Right-click any repository listed in the storage section, then select *Remove Repository*.  
The repository is removed from those listed for the selected host server and is no longer available to the host server.

### 3.12.2 Removing a Repository from the VM Client

---

**WARNING:** If a repository that you delete has VMs stored in it, they can no longer be managed.

---

- 1 In the **VM Client**, click the *Storage* view, right-click a repository, then select *Delete*.  
The repository is removed from those listed.

# Configuring Virtual Machines

# 4

You can configure existing virtual machines (VMs) in the PlateSpin® Orchestrate VM Client:

- ◆ [Section 4.1, “Editing VMs or Templates,” on page 45](#)
- ◆ [Section 4.2, “Moving VMs,” on page 53](#)
- ◆ [Section 4.3, “Migrating VMs,” on page 54](#)
- ◆ [Section 4.4, “Deleting VMs,” on page 56](#)
- ◆ [Section 4.5, “Deleting Templates,” on page 57](#)
- ◆ [Section 4.6, “Cloning Templates,” on page 57](#)
- ◆ [Section 4.7, “Detaching Clones from Templates,” on page 58](#)
- ◆ [Section 4.8, “Creating and Populating Groups,” on page 59](#)

## 4.1 Editing VMs or Templates

VMs or templates can be edited at any time. However, the fields and sections that are displayed on each of the four wizard pages depend on various circumstances:

- ◆ Which type of VM or template is it? Xen, VMware, or Hyper-V?
- ◆ What is the VM’s installed state? Defining, defined, installing, or installed?
- ◆ What is the VM’s defined state? Defining or defined?
- ◆ What is the run time state of the VM? Running, paused, suspended, or stopped?
- ◆ Is the VM state unknown?

Any combination of these factors determines which fields can be edited or shown. The displayed fields that can’t be edited are inaccessible (dimmed), or if no fields on a wizard page are editable, the page isn’t shown.

---

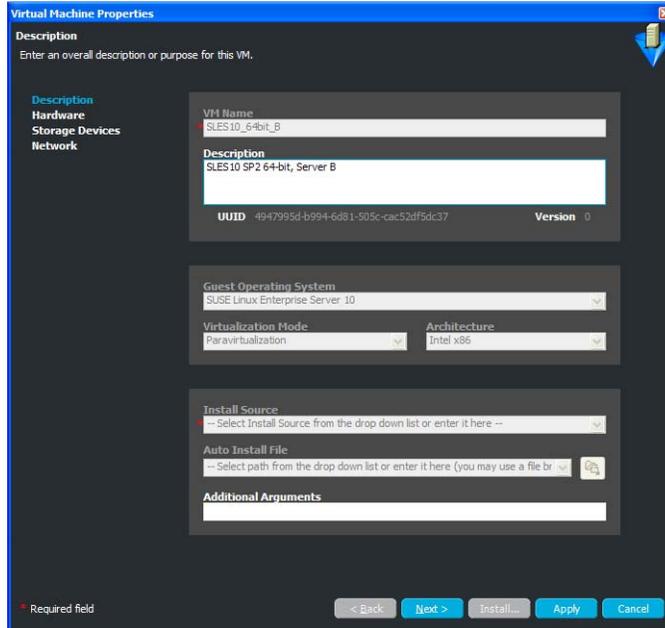
**IMPORTANT:** If the VM that you want to edit is running, only a few of its properties can be changed. If you want to edit a running VM, know that any changes are only retained in RAM, so when this instance if the VM is stopped, those edits are not kept.

---

To edit a VM or template, run the Virtual Machine Properties Wizard:

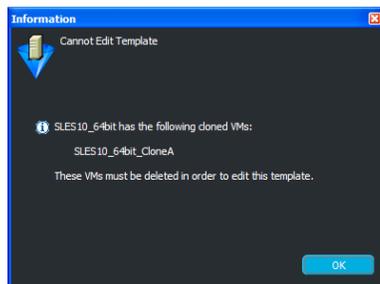
- 1 If you want to retain your edits on a VM, in the **VM Client**, click the *Virtual Machines* view, right-click the VM that you want to edit, then select *Shut Down*.
- 2 Do one of the following to start the wizard:
  - ◆ In the *Virtual Machines* view,
    - ◆ Select a VM, then click the **Edit** button.
    - ◆ Select a VM, then click *Edit > Edit VM*.
    - ◆ Select a VM, then press Ctrl+E.
    - ◆ Right-click any listed VM, then select *Edit*.

- ◆ In the *Templates* view,
  - ◆ Select a template, then click the **Edit** button.
  - ◆ Select a template, then click *Edit* > *Edit VM*.
  - ◆ Select a template, then press Ctrl+E.
  - ◆ Right-click any listed template, then select *Edit*.
- ◆ The following wizard page is displayed:



The example shown is for a stopped VM.

If a template currently has cloned VMs associated with it, you cannot edit the template. The following dialog box is displayed when you attempt to edit such a template:



After noting which VMs are involved, click *OK* to close the dialog box, then determine whether you want to detach those VMs in order to edit the template. For more information on detaching the clones from the template in the Development Client, see [Section 4.7, “Detaching Clones from Templates,” on page 58](#).

- 3 Fill in the available fields. For templates, most of the fields cannot be edited.
  - ◆ **VM Name:** (Required) Name for the VM or template.  
This name applies to all versions of the VM or template and must be unique. It cannot already exist in either the *Virtual Machines* or *Templates* listings.

The name cannot be longer than 100 characters and cannot start with the letters “xen.”

- ♦ **Description:** Note the purpose and any other specifics you might require of the VM or template. These notes apply to all versions of this VM or template.
- ♦ **UUID:** Displays the UUID for the VM or template. This field is view-only.
- ♦ **Guest Operating System:** The operating system you want to install on your VM or template.
- ♦ **Virtualization Mode:** *Full Virtualization* means that all the hardware components are emulated by the hypervisor. *Paravirtualization* means that the operating system directly accesses elements of the hardware, including processor, display, memory, network interface card, and hard drive, and uses the hypervisor to emulate other hardware interactions. Paravirtualization performs much faster and with better hardware efficiency than full virtualization.

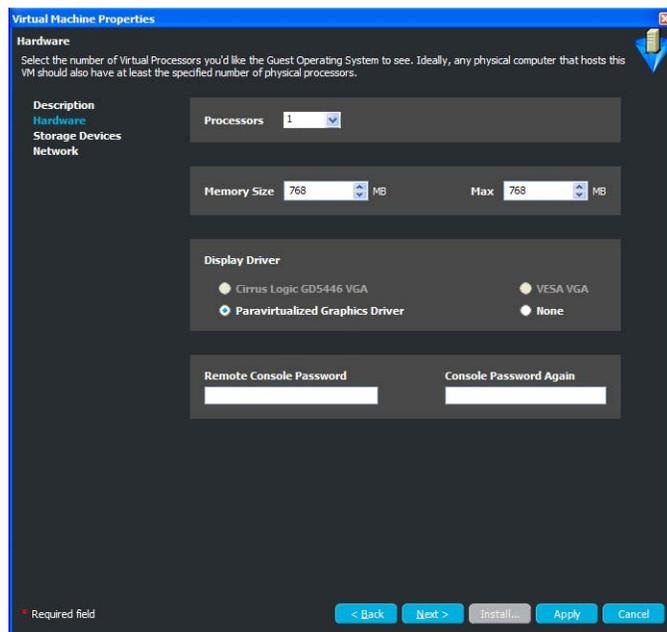
---

**NOTE:** Refer to the SUSE® Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL) documentation to verify supported platforms in each virtualization mode.

---

- ♦ **Architecture:** The type of processor you want to have as the virtual processor for the VM or template.
- ♦ **Install Source:** (Required) The path to the installation media for your guest operating system. For a fully virtualized VM, the install source must be an ISO. For a paravirtualized VM, an ISO or network install source is required. For more information, see [Appendix A, “VM Installation Sources,” on page 93](#).
- ♦ **Auto Install File:** This field applies only to paravirtual installations. The file listed here is used to populate the answers needed during the installation of the guest operating system (some examples include time zone, language, firewall, and other settings given in response to the user interactive areas of the installation). An AutoYaST file listed here can contain network settings that are applied during the VM build job. For RHEL, the file is a kickstart file.
- ♦ **Additional Arguments:** This field applies only to paravirtual installations. These are Linux install arguments, which are operating system parameters similar to the information entered in a GRUB boot loader when running on bare hardware in SLES and SUSE Linux Enterprise Desktop (SLED). Additional network installation information can be added here, which is used during the build process to find the installation source for the VM build process.

4 Click *Next* to display the following wizard page:

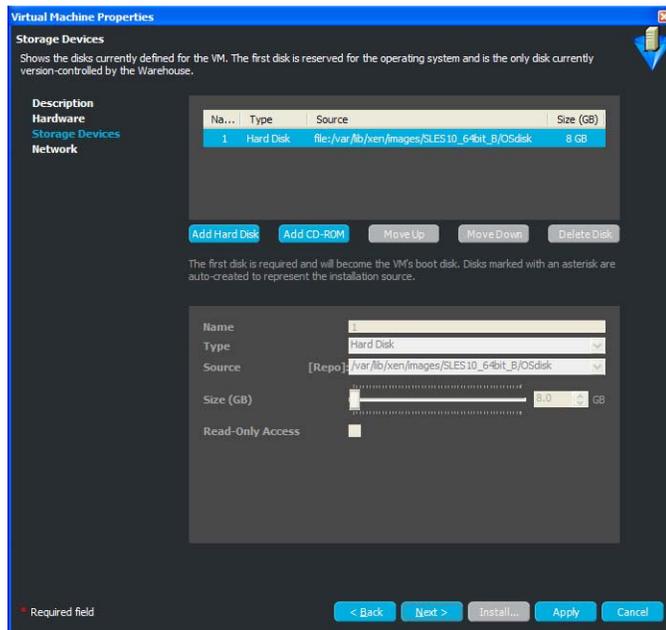


The example shown is for a stopped VM.

5 Fill in the available fields to define the kind of hardware to be virtualized for the VM or template:

- ◆ **Processors:** The number of virtual processors the VM or template has.
- ◆ **Memory Size:** The amount of virtual RAM the VM or template has, or paravirtualized memory that can be used by the VM if it is paravirtualized.
- ◆ **Max:** The amount of virtual RAM the VM or template has, or paravirtualized memory that can be used by the VM if it is paravirtualized.
- ◆ **Display Driver:** The properties of the VM's display. If you have chosen paravirtualization for this VM, your only option is *Paravirtualized*. If you have chosen full virtualization, you can select between *VESA* or the on-board VGA. If you do not need a display, you can turn off the display on a fully virtualized machine.
- ◆ **Remote Console Password:** The password for remotely controlling the VM. Use the two fields to confirm that the password is typed correctly.

6 Click *Next* to display the following wizard page:



The example shown is for a stopped VM.

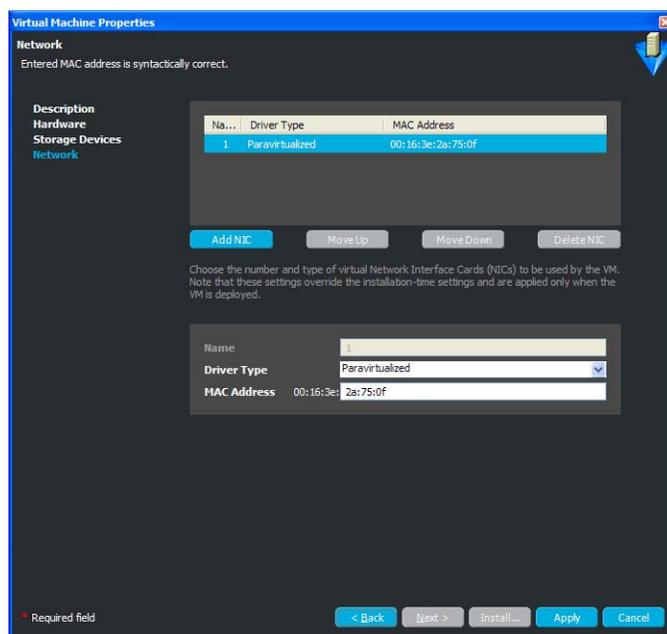
The Storage Devices page is for defining the hard drives and CD or DVD drives that you want to include in the VM or template profile. The first device must always be a hard disk that cannot be removed unless another hard disk has been put into the first slot. This is the only disk (Disk 0) that is stored and version controlled. All other disks are virtual and are created upon launching the VM. You can add any other storage devices you want to the VM or template and move them according to priority in the list.

7 To fill in the storage device information, select a task from the following table:

Tasks	Steps	Additional Details
Adding a hard disk to your VM	<ol style="list-style-type: none"> <li>Click <i>Add Hard Disk</i>.</li> <li>Fill in the following fields: <ul style="list-style-type: none"> <li><b>Name:</b> The name is the number in the order the drive appears in its VM image listing.</li> <li><b>Type:</b> Select the type of hard disk you want to create on the VM: <ul style="list-style-type: none"> <li>Hard Disk</li> <li>iSCSI</li> <li>N_Port ID Virtualized Storage</li> <li>Network Attached Storage</li> </ul> </li> <li><b>Source:</b> The repository on the hard disk of the VM.</li> <li><b>Size (GB):</b> The VM's hard drive size, ranging from 0.25 GB (250 MB) to 4 TB. This is the maximum allowable size of the VM.</li> <li><b>Read Only Access:</b> Select this option to give the hard drive Read-Only access. By default, hard drives are writable.</li> <li><b>Fabric ID:</b> (Available only for the N_Port ID virtualized storage). Select or specify a <b>fabric ID</b> that uniquely identifies a SAN repository. The <i>Fabric ID</i> drop-down list is populated with fabric IDs of the discovered SAN repositories.</li> <li><b>World Wide Name:</b> (Available only for the N_Port ID virtualized storage). Select or specify a unique <b>World Wide Name</b>. The World Wide Name value is automatically assigned for the NPIV disks. You can edit it, if required.</li> <li><b>Target ID:</b> (Available only for the N_Port ID virtualized storage). Specify the <b>target ID</b> of the SAN repository.</li> <li><b>LUN:</b> (Available only for the N_Port ID virtualized storage). Specify a <b>LUN</b> value between 0 and 254.</li> </ul> </li> <li>Click <i>Apply</i>.</li> </ol>	<p>The following restrictions apply per VM:</p> <ul style="list-style-type: none"> <li>Limit of 4 fully virtualized disks (hda-hdd).</li> <li>Limit of 16 paravirtualized disks (xvda-xvdp).</li> </ul> <p>While adding an additional NPIV disk to a VM, ensure that the disk values for Fabric ID, World Wide Name, Target ID, and LUN are unique for each NPIV disk of the VM.</p> <p>Creation of a block-based VM is not supported for NPIV.</p>

Tasks	Steps	Additional Details
Adding a CD-ROM or DVD ROM to your VM	<ol style="list-style-type: none"> <li>1. Click <i>Add CD-ROM</i>.</li> <li>2. Fill in the following fields: <ul style="list-style-type: none"> <li><b>Name:</b> The name is the number in the order the drive appears in its VM image listing.</li> <li><b>Type:</b> Select the type of CD-ROM you want to create on the VM: <ul style="list-style-type: none"> <li>◆ CD-ROM</li> <li>◆ DVD</li> </ul> </li> <li><b>Source:</b> The repository on the hard disk of the VM.</li> <li><b>Size (GB):</b> The size is zero and cannot be changed.</li> <li><b>Read Only Access:</b> These drives can only have Read-Only access.</li> </ul> </li> <li>3. Click <i>Apply</i>.</li> </ol>	
Rearranging the order in which each storage device is accessed	<ol style="list-style-type: none"> <li>1. Select the storage device whose order you want to change.</li> <li>2. Click <i>Move Up</i> or <i>Move Down</i>.</li> </ol>	The first storage device in the list must be a hard disk, which is the only one actually stored in the VM image.
Deleting a storage device	<ol style="list-style-type: none"> <li>1. Select the storage device you want to delete.</li> <li>2. Click <i>Delete</i>.</li> </ol>	Removes the selected storage device from the table.

8 Click *Next* to display the following wizard page:



The example shown is for a stopped VM.

The network devices are added, deleted, and defined on the Network page, where you add specific information about your network connections or intended network connections for your VM or template. You can define more detailed NIC settings in your auto-install files or when performing a guest installation.

**9** To fill in the NIC information, select a task from the following table:

Tasks	Steps
Adding a NIC to your VM	<ol style="list-style-type: none"> <li>1. Click <i>Add NIC</i>.</li> <li>2. Fill in the following fields: <ul style="list-style-type: none"> <li>◆ <b>Name:</b> Name of the NIC card. This name is used to help you to recognize each NIC.</li> <li>◆ <b>Driver Type:</b> Indicates whether the VM has direct access to the NIC. If the NIC should only interface with the hypervisor, you can select the specific type of full virtualization. You should have a paravirtualized NIC with a paravirtualized VM.</li> <li>◆ <b>MAC Address:</b> Unless you have a specific MAC address you want to assign to this VM, leave this set to <i>Default</i>, and Xen assigns it a MAC address when you launch the VM.</li> </ul> </li> <li>3. Click <i>Apply</i>.</li> </ol>
Rearranging the order in which the NICs are accessed	<ol style="list-style-type: none"> <li>1. Select the storage device whose order you want to change.</li> <li>2. Click <i>Move Up</i> or <i>Move Down</i>.</li> </ol>
Deleting a NIC from your VM	<ol style="list-style-type: none"> <li>1. Select the NIC you want to remove.</li> <li>2. Click <i>Delete NIC</i>.</li> </ol>

**10** To complete and close the wizard, click *Apply*.

You can click *Back* to make further changes before applying the edits.

Applies the wizard configuration changes and saves the VM in the *Virtual Machines* view, or the template in the *Templates* view.

Focus is returned to the *Virtual Machines* view after you exit the wizard with the edited VM or template selected in the list.

To install an edited VM at a later time, see [Section 3.6, “Installing a VM,” on page 38](#).

After editing a template, you can then use it to clone new VMs. For more information, see [Section 4.6, “Cloning Templates,” on page 57](#).

## 4.2 Moving VMs

When you need to change a VM's host server, you can either move or migrate the VM:

- ♦ **Moving:** For host resource reallocation. Physically moves all of the VM's files and its image from one host server to another while the VM is shut down. When you start the VM, it starts from its new repository.
- ♦ **Migrating:** For host server availability. Moves just the VM's instance in RAM from one host server to another (essentially keeping the VM live), assuming that both hosts have shared access to the VM's files and image. When you shut down the VM and later restart it, it restarts from the new host server. In other words, its files are not moved, but it now has a new host server association. For more information, see [Section 4.3, "Migrating VMs," on page 54](#).

Review the following sections to move a VM:

- ♦ [Section 4.2.1, "Prerequisites," on page 53](#)
- ♦ [Section 4.2.2, "Moving a VM," on page 53](#)

### 4.2.1 Prerequisites

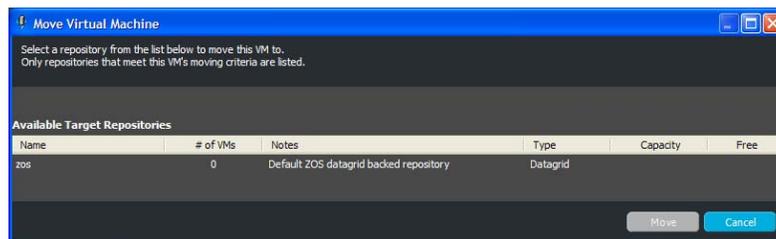
In order for a move to work, the following conditions must exist:

- ♦ The provisioning adapter for the VM being moved must match the repository's enabled provisioning adapters. (For this information, see the [Summary details tab](#) for the repository.)
- ♦ The target repository's device must have enough disk space available to hold the VM's files and image.

### 4.2.2 Moving a VM

- 1 In the **VM Client**, click the *Virtual Machines* view, then make sure the VM you want to move is shut down.
- 2 Select the VM.
- 3 Do one of the following to move the VM:
  - ♦ Click *Edit > Move*.
  - ♦ Right-click the selected VM, then select *Move*.

The following dialog box is displayed:



- 4 Under *Available Target Repositories*, select a repository.

The VM hosts connected to the repository are displayed. When you select the repository, you are selecting the host to move the VM to.

## 5 Click *Move*.

Depending on your hardware capabilities, moving the files can take a few minutes.

You can view the move progress in any of the following ways:

- ♦ Double-click the VM being moved, click the *Event Log* tab, then double-click the entry related to the moving process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
- ♦ Observe messages at the bottom of the VM Client interface.
- ♦ Click *View > Show Progress View* to open the Progress window.
- ♦ View its progress in the *Jobs* tab of the **Development Client**.

## 6 Start the VM from its new location.

# 4.3 Migrating VMs

When you need to change a VM's host server, you can either move or migrate the VM:

- ♦ **Migrating:** For host server availability. Moves just the VM's instance in RAM from one host server to another (essentially keeping the VM live), assuming that both hosts have shared access to the VM's files and image. When you shut down the VM and later restart it, it restarts from the new host server. In other words, its files are not moved, but it now has a new host server association.
- ♦ **Moving:** For host resource reallocation. Physically moves all of the VM's files and its image from one host server to another while the VM is shut down. When you start the VM, it starts from its new repository. For more information, see [Section 4.2, "Moving VMs," on page 53](#).

Review the following sections to migrate a VM:

- ♦ [Section 4.3.1, "Prerequisites," on page 54](#)
- ♦ [Section 4.3.2, "Migrating a VM," on page 55](#)

## 4.3.1 Prerequisites

In order for a migration to work, the following conditions must exist:

- ♦ Xen must be configured to allow migrations.
- ♦ The two host servers involved in the migration must have access to the shared repository where the VM's files and image reside.
- ♦ Every entity involved must have the same architecture, such as 32-bit or 64-bit, mount points, and routing of network connections to the virtual network.

For example, a 64-bit paravirtualized guest created on a 64-bit hypervisor can be migrated to another host server running a 64-bit hypervisor, but not a 32-bit hypervisor. Or, you can create a 32-bit VM on a 64-bit hypervisor and migrate it to another 64-bit hypervisor host. In other words, you cannot start a VM on a host that has a different architecture from the host where the VM was created.

- ♦ The host server cannot be fully utilized already, because it would be incapable of hosting the VM.

- ◆ The host server cannot have met its maximum allowable VMs already. This value is set in the Development Client in the *Max Hosted VMs* field. The default is 3. Also, it doesn't matter whether the maximum VMs it is already hosting are fully utilizing it.
- ◆ A host server cannot be disabled by having its *Resource Enabled* check box deselected in the Development Client. By default, this check box is enabled.
- ◆ A host server is only available if it is not out of sync, which can occur if it is manually accessed or manipulated via a third-party tool. However, you can **resynchronize** the host server with the Orchestrate Server.

These conditions are automatically enforced with **constraints**.

During migration:

- ◆ If the above criteria are met by only one host server, the VM is automatically migrated.
- ◆ If the criteria are met by multiple host servers, you can either select a host server, or to select to have the VM Client automatically select the host server.
- ◆ If no host servers meet the requirements, the No Migrate Target Available dialog box is displayed, indicating that you need to resolve the issues to continue.

### 4.3.2 Migrating a VM

To migrate a VM, you need to prepare the host servers, then migrate the VM.

- 1 Make sure that each host has access to the shared repository to be used for the VM's files.
- 2 In the **VM Client**, click the *Virtual Machines* view, select the VM to migrate, then do one of the following:
  - ◆ Click the **Migrate** button.
  - ◆ Click *Edit > Migrate*.
  - ◆ In the *Virtual Machines* view, right-click the selected VM, then select *Migrate*.

If only one host server is available, it is automatically used and you can continue with **Step 4**. Otherwise, the following dialog box is displayed if there are two or more target host candidates:



- 3 If you have multiple host servers to select from, do one of the following:
  - ◆ Leave the *Auto-select the Target Host* check box selected.

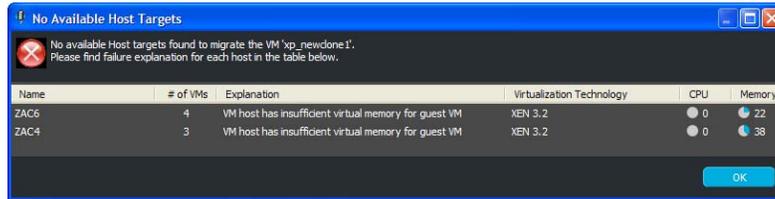
The *Auto-Select the Target Host* option allows the Orchestrate Server to automatically select a host server from those that are available by using a ranking criteria, such as architecture similarity and available CPU and RAM resources. For example, hosts are selected that meet the architecture requirements, the cost of moving information is then

considered, and if there is still no clear winner, then the least loaded machine is selected. The ranking criteria can be affected by policies that you might have set in the Development Client.

- ◆ Select the VM host to which you want to migrate the VM.

These options are mutually exclusive, meaning that selecting one disables the other.

If you do not yet have a valid host server to migrate to, the following dialog box is displayed:



Review the issues, resolve them, then repeat from [Step 1](#).

#### 4 Click *Migrate*.

You can view the migration progress in any of the following ways:

- ◆ Double-click the VM being migrated, click the *Event Log* tab, then double-click the entry related to the migration process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
- ◆ Observe messages at the bottom of the VM Client interface.
- ◆ Click *View > Show Progress View* to open the Progress window.
- ◆ View its progress in the *Jobs* tab of the [Development Client](#).

When the migration has completed, the VM continues to run on its new host.

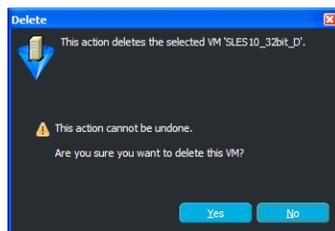
## 4.4 Deleting VMs

Deleted VMs cannot be recovered. Deleting a VM in the VM Client causes an Orchestrate job to run that deletes all of the VM's files in the `/etc/xen/vm` and `/etc/xen/images` (or `/var/lib/xen/images`) directories.

- 1 In the [VM Client](#), click the *Virtual Machines* view, then select one or more VMs.
- 2 Do one of the following to delete the selected VMs:
  - ◆ Click the  button.
  - ◆ Click *Edit > Delete*.
  - ◆ In the *Virtual Machines* view, right-click a selected VM, then select *Delete*.

If a VM is running or is suspended, the deletion options are not displayed or enabled.

- 3 Click either *Yes* or *No* to delete the VM:

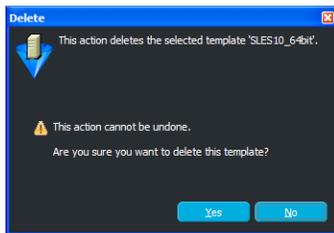


This completely removes the selected VMs from the *Virtual Machines* view, as well as all files in the repository associated with it. Deleted VMs cannot be recovered.

## 4.5 Deleting Templates

Deleted templates cannot be recovered. You also cannot delete a template that has clones. However, you can detach the clones from the template in the Development Client. For more information, see [Section 4.7, “Detaching Clones from Templates,” on page 58](#).

- 1 In the **VM Client**, click the *Templates* view, then select one or more templates.
- 2 Do one of the following to delete the selected templates:
  - ♦ Click the **Delete** button.
  - ♦ Click *Edit > Delete*.
  - ♦ In the *Templates* view, right-click the selected template, then select *Delete*.
- 3 Click either *Yes* or *No* to delete the template:



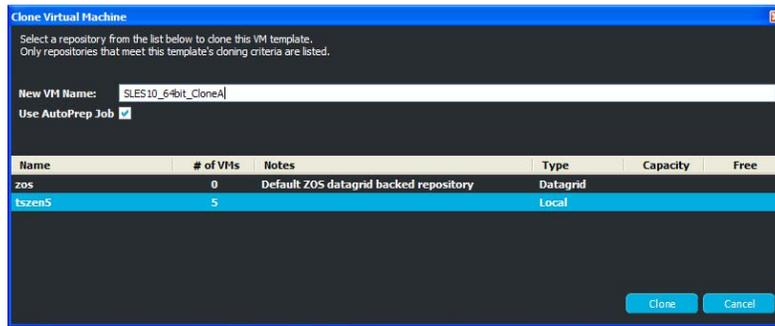
This completely removes the templates from the *Templates* view. It does not have any affect on any VMs cloned from these templates.

## 4.6 Cloning Templates

Templates provide a method for quickly creating of a new VM. To create new VMs from a template, you first clone the template, then you edit the new VM’s configuration as needed.

- 1 In the **VM Client**, click the *Templates* view, then select a template.
- 2 Do one of the following to clone the template:
  - ♦ Click the **Clone** button.
  - ♦ Click *Edit > Clone*.
  - ♦ In the *Templates* view, right-click the selected template, then select *Clone*.

The following dialog box opens:



- 3 Specify a name for the template clone in the *New VM Name* field.

This name applies to all versions of the clone and must be unique. It cannot already exist in either the *Virtual Machines* or *Templates* listings.

The name cannot be longer than 100 characters and cannot start with the letters “xen.”

- 4 To select a repository for the new template clone, simply click the repository.

This enables the *Clone* button.

The default repository that is displayed is the location where the template being cloned resides.

- 5 Select the host for the template.
- 6 Click *Clone*.

A cloning job is started. You can view its progress in any of the following ways:

- ♦ Double-click the template being cloned, click the *Event Log* tab, then double-click the entry related to the cloning process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
- ♦ Observe messages at the bottom of the VM Client interface.
- ♦ Click *View > Show Progress View* to open the Progress window.
- ♦ View its progress in the *Jobs* tab of the **Development Client**.

When the cloning process is complete, the VM Client’s job status is displayed in a pop-up message, as well as in the Event Log Details dialog box. It is also listed in the Development Client as *Completed*.

The cloned template is now listed with the other VMs in the *Virtual Machines* view.

Except for the name (see **Step 3**), the UUID, and the MAC address, the new template is identical to the template that you started with in **Step 1**.

- 7 (Optional) Right-click the cloned template, then select *Edit* to modify it as necessary.

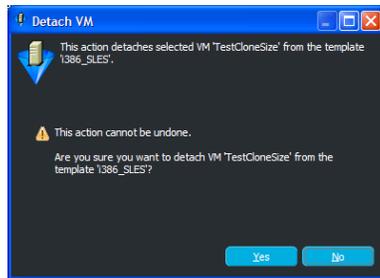
## 4.7 Detaching Clones from Templates

If you need to detach one or more cloned VMs in order to edit or delete a template, do the following:

- 1 In the **VM Client**, view the dialog box that is displayed when you attempt to edit or delete a template and note which VMs need to be detached.
- 2 Select *Virtual Machines* in the Inventory view, then do one of the following:
  - ♦ Select a VM to be detached, then click the **Detach** button.

- ◆ Select a VM to be detached, then click *Edit > Detach*.
- ◆ Right-click a VM to be detached, then select *Detach*.

The following dialog box is displayed:



- 3 Click *OK* to confirm the detachment of the VM from the template.

The VM is repositioned under *VMs* now that it is no longer related to a template.

- 4 For other VM clones, repeat [Step 2](#) through [Step 3](#).
- 5 Return to the VM Client to either [edit the template](#) or [delete the template](#).

## 4.8 Creating and Populating Groups

You can create groups in each of the four Inventory perspective's views for managing their listed items:

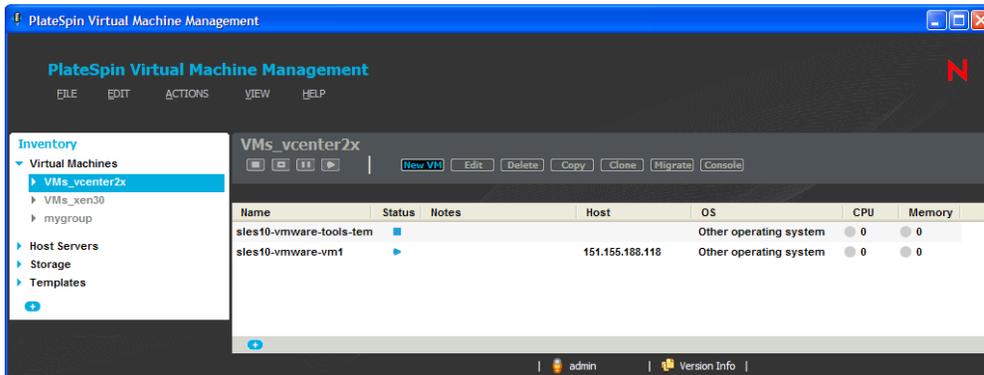
- ◆ [Section 4.8.1, "Understanding Groups," on page 59](#)
- ◆ [Section 4.8.2, "Creating a Group," on page 60](#)
- ◆ [Section 4.8.3, "Adding Members to an Existing Group," on page 60](#)
- ◆ [Section 4.8.4, "Removing Members from a Group," on page 61](#)
- ◆ [Section 4.8.5, "Deleting a Group," on page 61](#)

### 4.8.1 Understanding Groups

Groups allow you to better manage a view's items when you have too many items to view on one screen. You can create multiple groups under any of the four Inventory perspective views (Virtual Machines, Host Servers, Storage, and Templates).

The following figure illustrates groups in the Inventory perspective panel:

**Figure 4-1** Groups in the Inventory Perspective of the VM Client



## 4.8.2 Creating a Group

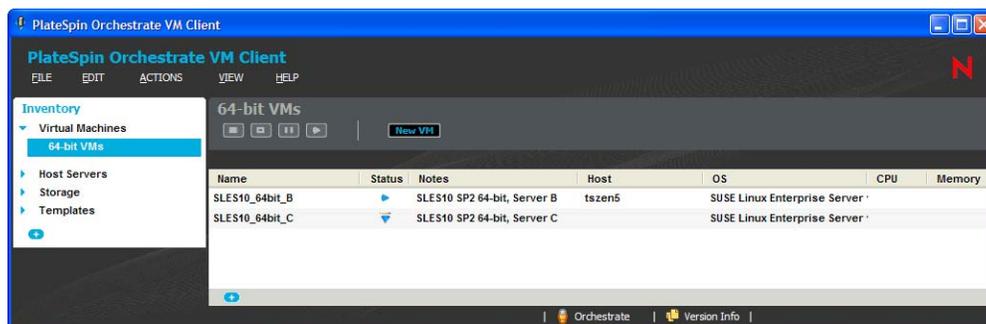
- 1 In the **VM Client**, do one of the following:
  - ◆ Right-click a view in the Inventory perspective, then click *New Group*.
  - ◆ Select a view in the Inventory perspective, then click *File > New Group*.

This opens the New Group dialog box:



- 2 Specify a group name, then click *OK*.

The group is now displayed under the selected view:



- 3 To add members to the group, continue with [Section 4.8.3, “Adding Members to an Existing Group,”](#) on page 60.

## 4.8.3 Adding Members to an Existing Group

- 1 In the **VM Client**, click a view in the Inventory perspective.

- 2 Click the view's  icon to expand its groups.
- 3 In the details panel, right-click one of the listed items (a VM, a host server, a repository, or a template), then drag and drop it into one of the view's existing groups.

The item is still listed in the view if you do not have a group selected for the view. That way, you can assign an item membership in multiple groups.

#### 4.8.4 Removing Members from a Group

- 1 In the **VM Client**, click a view in the Inventory perspective (Virtual Machines, Host Servers, Storage, or Templates).
- 2 Click the view's  icon to expand its groups, then select one of the groups.  
The details panel on the right now displays only the members of the group.
- 3 In the details panel, right-click one of the listed items (a VM, a host server, a repository, or a template).
- 4 Click *Remove from Group*.

The selected item is removed from the view's group, but not from the view's full listing.

#### 4.8.5 Deleting a Group

- 1 In the Inventory perspective in **the VM Client**, click a view's  icon to expand the view.
- 2 Right-click the group that you want to delete, then select *Delete*.

The group is removed from the Inventory perspective panel. However, the members of the group are not removed from the view's listing. They are only disassociated from the deleted group.



# Managing Virtual Machines

# 5

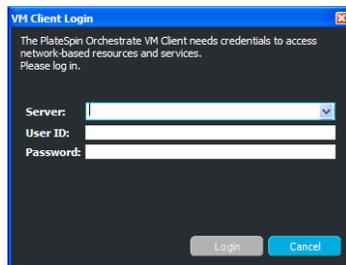
You can do the following for hands-on management of virtual machines (VMs) in the PlateSpin® Orchestrator VM Client:

- ◆ Section 5.1, “Logging In to a Virtualization Grid,” on page 63
- ◆ Section 5.2, “Understanding a VM’s Status,” on page 64
- ◆ Section 5.3, “Canceling a VM’s In-Progress Action,” on page 65
- ◆ Section 5.4, “Using Feedback to Manage VMs,” on page 65
- ◆ Section 5.5, “Using Feedback to Manage Host Servers,” on page 67
- ◆ Section 5.6, “Starting a VM,” on page 67
- ◆ Section 5.7, “Shutting Down a VM,” on page 69
- ◆ Section 5.8, “Accessing a VM’s Console,” on page 70
- ◆ Section 5.9, “Pausing a VM,” on page 70
- ◆ Section 5.10, “Resuming a Paused VM,” on page 70
- ◆ Section 5.11, “Suspending a VM,” on page 71
- ◆ Section 5.12, “Resuming a Suspended VM,” on page 71
- ◆ Section 5.13, “Resynchronizing a Host Server’s State,” on page 72
- ◆ Section 5.14, “Resynchronizing a VM’s State,” on page 72
- ◆ Section 5.15, “Using Groups to Manage Your VMs,” on page 73
- ◆ Section 5.16, “Using the Error Log View,” on page 73
- ◆ Section 5.17, “Using the Progress View,” on page 73
- ◆ Section 5.18, “Resetting the VM Client Perspectives,” on page 74
- ◆ Section 5.19, “Viewing and Editing VM Client Details,” on page 74

## 5.1 Logging In to a Virtualization Grid

You can log in to a different Orchestrator Server at any time for access to a virtualization grid. However, you cannot be logged in to multiple Orchestrator Servers. If you log in to another server, the previous login is disabled.

- 1 In the **VM Client**, click *File > Log In*.



If you are already logged in to a virtualization grid, you can press Ctrl+Shift+I to log in to a different grid.

**2** Fill in the fields:

- ♦ **Server:** Specify either the DNS name or IP address of the Orchestrate Server.
- ♦ **User ID:** Specify a username that has admin rights on the server.
- ♦ **Password:** Specify this user’s password.

**3** Click *Login*.

## 5.2 Understanding a VM’s Status

In the *Virtual Machines* view, the *Status* column indicates the state of each VM:

**Table 5-1** VM Status Icons

Icon	VM’s Status	Enabled Actions for the Selected VM
	Has been defined or is being defined.	<i>Show Details, Install, Edit, Resync State, Delete, New VM, and Remove from Group.</i>  The VM is eligible to be installed.
	Has an action that is in progress.	Other actions cannot be run against this VM until the current action has finished.
	Is running.	<i>Show Details, Restart, Suspend, Pause, Shut Down, Edit, Resync State, Migrate, New VM, Show Console, and Remove from Group.</i>
	Is stopped.	<i>Show Details, Start, Edit, Create Template, Resync State, Move, Delete, Install Agent, New VM, and Remove from Group.</i>
	Is paused.	<i>Show Details, Resume, Edit, Resync State, Delete, New VM, and Remove from Group.</i>
	Is suspended.	<i>Show Details, Resume, Edit, Resync State, Move, Delete, New VM, and Remove from Group.</i>
	The state is unknown.	<i>Show Details, Edit, Resync State, Delete, New VM, and Remove from Group.</i>  For information on making an unknown state known, see <a href="#">Section 5.13, “Resynchronizing a Host Server’s State,”</a> on page 72 and <a href="#">Section 5.14, “Resynchronizing a VM’s State,”</a> on page 72.

VM statuses are updated globally; they show any status change that is made in your local instance of the VM Client and any change made from other VM Client instances being run by other administrators.

You can tell whether a VM or host server has the Orchestrate Agent installed on it by whether there are pie charts in its *CPU* and *Memory* columns.

## 5.3 Canceling a VM's In-Progress Action

Sometimes the progress of an action taken on a VM gets hung up, or you might change your mind on whether to proceed with the action. You can do the following:

- ♦ To determine whether an action is progressing normally, double-click the  icon to open the *Even Log* tab for the VM and review the log entries.

For more information, see [Section 5.4, “Using Feedback to Manage VMs,” on page 65](#).

- ♦ To cancel an action, go to the *Explorer* tree view in the **Development Client**. Under the *Resources* subsection, locate the VM in one of the *Resources*' subsections (such as *VMs*), right-click the VM, then select *Cancel Action*.

In the VM Client window, the *Status* column should now show the previous state of the VM.

## 5.4 Using Feedback to Manage VMs

The VM Client is integrated with the Orchestrate Server and the **Development Client** to provide details from any Orchestrate jobs that are started by actions you take in the VM Client. This provisioning log integration enables you to see the current status on jobs that are running and to review the status of jobs that have completed. Therefore, you do not need to open the Development Client to discover the status of jobs started in the VM Client.

The following table describes the various things you can do in the VM Client to see the progress and results of any actions that you take:

**Table 5-2** *Ways to View Actions' Progress and Results*

Action	Keystrokes	Additional Information
Open the Progress view for a VM that is in a progress or unknown state	Click <i>View &gt; Show Progress View</i>	The  icon is displayed in the <i>Status</i> column when a VM is in a progress state.
	Click the  icon to the right of the green progress bar in the lower right of the VM Client window	The Progress view displays progress bars for actions taken globally, meaning any taken in your local instance of the VM Client, as well as from any other instance of the VM Client running for current host server.
	Double-click the text to the left of the green progress bar in the lower right of the VM Client window	

Action	Keystrokes	Additional Information
View a progress bar for an action taken on a VM	Turn your focus to the upper right of the VM Client window	<p>An area in the upper right of the VM Client window shows both a progress bar and a short description of the action:</p>  <p>Click the bar or the description to access the <i>Event Log</i> tab for the action in progress.</p> <p>This progress bar is displayed only for the last action taken in the local VM Client.</p>
	Turn your focus to the lower right of the VM Client window	<p>An area in the lower right of the VM Client window shows both a progress bar and a short description of the action:</p>  <p>Click the bar or the description to access the <i>Event Log</i> tab for the action in progress.</p> <p>This progress bar is displayed only for the last action taken by any VM Client that is open, meaning that this is a global progress bar.</p>
View a detailed tooltip of the status of a VM that is in a progress state	Mouse over the VM's ✱ icon	<p>The ✱ icon is displayed in the <i>Status</i> column when a VM is in a progress state.</p> <p>While the VM is in a progress state, other actions cannot be performed on the VM because <b>Orchestrate jobs</b> are run sequentially per VM.</p>
View the <i>Event Log</i> tab for a VM that is in a progress or unknown state	Double-click the VM in the <i>Virtual Machines</i> listing	This automatically opens the <i>Event Log</i> tab in the details view for the VM, instead of opening the <i>Summary</i> tab (the default double-click action).
	Single-click the description next to the progress bar located in the upper right of the VM Client window	
	Double-click the progress bar located in the upper right of the VM Client window	
	Double-click the entry in the Progress view (if it is open)	

Action	Keystrokes	Additional Information
View an <i>Event Log</i> entry's job details	Double-click the  icon on an <i>Event Log</i> tab's entry	<p>Provides further job details for that <i>Event Log</i> entry. These details are logged dynamically in the Event Log Details dialog box.</p> <p>Only the <i>Event Log</i> entries that have further job information available have the  icon in the <i>Additional Information</i> column.</p> <p>You can have multiple Event Log Details dialog boxes open simultaneously to monitor multiple tasks for the same VM or for multiple VMs.</p> <p>The Event Log Details dialog box provides the same information as the Job Details dialog box in the Development Client.</p>

## 5.5 Using Feedback to Manage Host Servers

The VM Client is integrated with the Orchestrate Server and the **Development Client** to provide event details for the host servers that are listed in the VM Client. This integration with the Orchestrate Server resource log enables you to see details on any jobs run on each host. Therefore, you do not need to open the Development Client to discover the events for the host servers that are listed in the VM Client.

The *Event Log* tab for a host server shows all jobs that have been run for VMs hosted by that server. This allows you to view all provisioning and life cycle actions that are being or have been run on that host server.

Some actions for a host server that aren't yet associated with a resource in the system, such as `Accept Pending Registration` for the host, can be reviewed for that host after the action has been completed. You can then click any of the  icons in the *Additional Information* column of the *Event Log* tab for that host to obtain the details of any registration event. For example, the `Accept Pending Registration` action runs multiple jobs on the Orchestrate Server, causing multiple events to be logged. Each of those job-related entries in the *Event Log* tab might have further information, as indicated by the  icon.

## 5.6 Starting a VM

In order start a VM on a host server, the following conditions must exist:

- ◆ The VM's files and image must be either:
  - ◆ Stored locally on the host server.
  - or
  - ◆ Stored in a shared repository accessible by the host server.
- ◆ The host server and VM must have the same architecture, such as 32-bit or 64-bit, mount points, and routing of network connections to the virtual network.

For example, a 64-bit paravirtualized guest created on a 64-bit hypervisor can be started on a host server running a 64-bit hypervisor, but not a 32-bit hypervisor. Or, a 32-bit VM created on a 64-bit hypervisor can be started on a 64-bit hypervisor host. In other words, you cannot start a VM on a host that has a different architecture from the host where the VM was created.

- ◆ The host server cannot be fully utilized already, because it would be incapable of hosting the VM.
- ◆ The host server cannot have met its maximum allowable VMs already. This value is set in the Development Client in the *Max Hosted VMs* field. The default is 3. Also, it doesn't matter whether the maximum VMs it is already hosting are fully utilizing it.
- ◆ The VM cannot be disabled by having its *Resource Enabled* check box deselected in the Development Client. By default, this check box is enabled.
- ◆ The VM cannot be out of sync, which can occur if it is manually accessed or manipulated via a third-party tool. For example, you could manually start the VM on the physical host machine. However, you can **resynchronize** the VM with the Orchestrate Server.

Host servers are only available if they are not **out of sync**, as well.

These conditions are automatically enforced with **constraints**.

If the above criteria are met by the VM and only one host server, then the VM is automatically started.

If the criteria are met by multiple host servers, you can start the VM by selecting a host server, or by selecting to have the Orchestrate Server automatically select the host server.

If no host servers meet these requirements, or if the VM doesn't meet its requirements, the No Start Target Available dialog box is displayed, indicating that you need to resolve the issues to continue.

You can start any VM that is not running. However, starting a paused or suspended VM can cause its previously saved state to be lost. For more information on paused or suspended VMs, see [Section 5.9, "Pausing a VM," on page 70](#) and [Section 5.11, "Suspending a VM," on page 71](#).

To start a VM:

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Do one of the following:
  - ◆ Click the VM that you want to start, then click the  button.
  - ◆ Click the VM that you want to start, then click *Action > Start*.
  - ◆ Right-click the VM that you want to start, then select *Start*.

If only one host server is available, it is automatically used. If there are two or more target host candidates, the following dialog box is displayed:



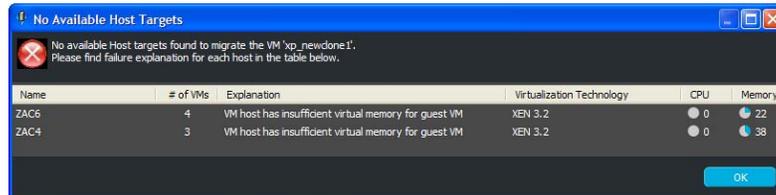
- 3 Select a host server, or select the *Auto-Select the Target Host* check box.

The *Auto-Select the Target Host* option allows the Orchestrate Server to automatically select a host server from those that are available by using a ranking criteria, such as architecture similarity and available CPU and RAM resources. For example, hosts are selected that meet the

architecture requirements, the cost of moving information is then considered, and if there is still no clear winner, then the least loaded machine is selected. The ranking criteria can be affected by policies that you might have set in the Development Client.

The VM starts, using the processes for the operating system it is running. You can use the *Show Console* menu option to authenticate if you want to view its console and log in to the operating system.

If there is no available host server for starting the VM, the following dialog box is displayed:



- 4 Review the issues, resolve them, then repeat from [Step 2](#).
- 5 You can view the startup progress in any of the following ways:
  - ♦ Double-click the VM being started, click the *Event Log* tab, then double-click the entry related to the starting process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
  - ♦ Observe messages at the bottom of the VM Client interface.
  - ♦ Click *View > Show Progress View* to open the Progress window.
  - ♦ View its progress in the *Jobs* tab of the [Development Client](#).

## 5.7 Shutting Down a VM

Shutting down a VM is the same as shutting down any computer. Its current state is not saved. If you want to retain a VM's state, use one of the following options instead:

- ♦ [Section 5.9, "Pausing a VM," on page 70](#)
- ♦ [Section 5.11, "Suspending a VM," on page 71](#)

To shut down a VM:

- 1 In the [VM Client](#), click the *Virtual Machines* view.
- 2 Do one of the following:
  - ♦ Select the VMs to be shut down, then click the  button.
  - ♦ Select the VMs to be shut down, then click *Actions > Shut Down*.
  - ♦ Right-click each VM to be shut down, then select *Shut Down*.

The selected VMs are shut down. Their current state is not saved.

- 3 You can view the shutdown progress in any of the following ways:
  - ♦ Double-click the VM being shut down, click the *Event Log* tab, then double-click the entry related to the process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
  - ♦ Observe messages at the bottom of the VM Client interface.

- ♦ Click *View > Show Progress View* to open the Progress window.
  - ♦ View its progress in the *Jobs* tab of the **Development Client**.
- 4 To restart a VM that is shut down, continue with [Section 5.6, “Starting a VM,”](#) on page 67.

## 5.8 Accessing a VM’s Console

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Do one of the following:
  - ♦ Select the VM, then click the  button.
  - ♦ Select the VM, then click *View > Show Console*.
  - ♦ Right-click the VM, then select *Show Console*.
- 3 Provide the VNC password, then click *OK*.

## 5.9 Pausing a VM

If you need to free up CPU cycles on the host machine, you can pause some of the VMs. This option writes the current state of the VM to the host machine’s RAM so that it can be recovered by [resuming the VM](#). A paused VM can also be moved to a suspended state without resuming it first.

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Do one of the following:
  - ♦ Select the VMs to be paused, then click the  button.
  - ♦ Select the VMs to be paused, then click *Actions > Pause*.
  - ♦ Right-click the VMs to be paused, then select *Pause*.

Each VM’s state is written to RAM on the host machine.

---

**WARNING:** If you use the *Restart* option to reboot a VM, its previous state is lost. If you want to recover its previous state, use the [Resume option](#), which reads the VM information from the host machine’s RAM.

Also, if you reboot the host machine, the VM’s state is lost, because it was saved to the host machine’s RAM. To prevent the loss of a VM’s state when you need to reboot the host machine, first use the [Suspend option](#) to suspend the VM.

---

## 5.10 Resuming a Paused VM

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Do one of the following:
  - ♦ Select the VMs to start again, then click the  button.
  - ♦ Select the VMs to start again, then click *Actions > Resume*.
  - ♦ Right-click the VMs to start again, then select *Resume*.

The previous state of the VM is restored from the information that was written to the host machine’s RAM when the *Pause* option was used.

If the VM’s console was open, it reopens automatically.

## 5.11 Suspending a VM

If you need to temporarily free up RAM on your host machine, or to save the state of any VMs running when you reboot the host machine, use the *Suspend* menu option. This option writes the current state of the VM to the host machine's hard drive so that it can be recovered by **resuming the VM**.

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Do one of the following:
  - ♦ Select the VMs to be suspended, then click the  button.
  - ♦ Select the VMs to be suspended, then click *Actions > Suspend*.
  - ♦ Right-click the VMs to be suspended, then select *Suspend*.

When you select multiple VMs, each suspend operation is done according to how **jobs** are handled by Orchestrate Server.

---

**WARNING:** If you use the *Restart* option to reboot the VM, its previous state is lost. If you want to recover its previous state, use the **Resume option**, which reads that information from the host machine's hard drive.

---

- 3 You can view the progress in any of the following ways:
  - ♦ Double-click the VM being suspended, click the *Event Log* tab, then double-click the entry related to the process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
  - ♦ Observe messages at the bottom of the VM Client interface.
  - ♦ Click *View > Show Progress View* to open the Progress window.
  - ♦ View its progress in the *Jobs* tab of the **Development Client**.

## 5.12 Resuming a Suspended VM

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Do one of the following:
  - ♦ Select the VMs to start again, then click the  button.
  - ♦ Select the VMs to start again, then click *Actions > Resume*.
  - ♦ Right-click the VMs to start again, then select *Resume*.

The previous state of the VM is restored from the information that was written to the host machine's hard drive when the *Suspend* option was used.

If you had the VM's console open, you need to reopen it by using the *Show Console* menu option.

- 3 You can view the progress in any of the following ways:
  - ♦ Double-click the VM being resumed, click the *Event Log* tab, then double-click the entry related to the process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
  - ♦ Observe messages at the bottom of the VM Client interface.
  - ♦ Click *View > Show Progress View* to open the Progress window.
  - ♦ View its progress in the *Jobs* tab of the **Development Client**.

## 5.13 Resynchronizing a Host Server's State

The Orchestrate Server is automatically updated to list the state of the hosted VMs after every **provisioning** action. However, the state of the hosted VMs is not automatically updated if the VM state is changed by any third-party tool. You must manually synchronize the state of the hosted VMs in the Orchestrate Server to reflect the current state of the VMs.

- 1 In the **VM Client**, click the *Host Servers* view.
- 2 Right-click the VM host to be resynchronized, then select *Resync State*.

This resynchronizes the state of all VMs hosted by the VM host.

You cannot select multiple hosts for this operation.

If the resync state is not automatically updated for any of the hosted VMs, open the details page of the appropriate VM in any of the following ways to get the updated status.

- ♦ Double-click the VM.
  - ♦ Right-click the VM, then select *Show Details*.
- 3 You can view the progress in any of the following ways:
    - ♦ Double-click the host being synchronized, click the *Event Log* tab, then double-click the entry related to the process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
    - ♦ Observe messages at the bottom of the VM Client interface.
    - ♦ Click *View > Show Progress View* to open the Progress window.
    - ♦ View its progress in the *Jobs* tab of the **Development Client**.

## 5.14 Resynchronizing a VM's State

The Orchestrate Server is automatically updated on the VM's state after every **provisioning** action. However, the VM's state is not automatically updated in the Orchestrate Server if the VM state is changed by any third-party tool or by manually manipulating the VM on the physical host machine. You must manually synchronize the VM's state in the Orchestrate Server to reflect the VM's current state.

- 1 In the **VM Client**, click the *Virtual Machines* view.
- 2 Right-click the VM to be resynchronized, then select *Resync State*.

You cannot select multiple VMs for this operation.

If the resync state is not automatically updated for the VM, open the VM's details page for the VM in any of the following ways to get the updated status.

- ♦ Double-click the VM.
  - ♦ Right-click the VMs, then click *Show Details*.
- 3 You can view the progress in any of the following ways:
    - ♦ Double-click the VM being synchronized, click the *Event Log* tab, then double-click the entry related to the process that has the  icon next to it. The Event Log Details dialog box is displayed and automatically updated as events occur.
    - ♦ Observe messages at the bottom of the VM Client interface.

- Click *View > Show Progress View* to open the Progress window.
- View its progress in the *Jobs* tab of the **Development Client**.

## 5.15 Using Groups to Manage Your VMs

- 1 In the **VM Client**, click the  icon for an Inventory's view (*Virtual Machines, Host Servers, Storage, or Templates*) to expand its existing groups.

The icon is not displayed for sections that do not contain at least one group.

- 2 To view a group's membership, click the group's name in the Inventory view.

The group's members are displayed in the details panel on the right. You can work with those members the same way you work with the items in the view's full listing.

You can expand as many groups as you like. A vertical scroll bar provides navigation.

For more information on groups, see [Section 4.8, "Creating and Populating Groups,"](#) on page 59.

## 5.16 Using the Error Log View

The Error Log view is displayed at the bottom of the interface window when you click *View > Show Error Log*. You can also double-click an error message in the lower left of the VM Client window to open this view.

Both the Progress and Error Log views use the same position in the VM Client window. You can either drag them to other positions, or toggle between them by using the *View* menu options, or the  icon located in the lower right of the VM Client window. For more information, see [Section 2.5, "Understanding the VM Client View and Details Editors,"](#) on page 25.

You can double-click in the Error Log view's title area to expand it to use up the entire VM Client window. Double-click again to return it to its previous size and location.

The VM Client is based on Eclipse and uses the standard Eclipse Error Log view. For more information, see [Error Log \(http://help.eclipse.org/ganymede/topic/org.eclipse.pde.doc.user/guide/tools/views/error\\_log.htm\)](http://help.eclipse.org/ganymede/topic/org.eclipse.pde.doc.user/guide/tools/views/error_log.htm).

## 5.17 Using the Progress View

The Progress view is useful for when you have actions occurring simultaneously for multiple VMs because you can keep track of which ones are still running. The Progress view is displayed at the bottom of the interface window when you click *View > Show Progress View*. You can also double-click a progress message in the lower left of the VM Client window to open this view.

**Figure 5-1** *Progress View for the VM Client*



Both the Progress view and Error Log views use the same position in the VM Client window. You can either drag them to other positions, or toggle between them by using the *View* menu options. For more information, see [Section 2.5, “Understanding the VM Client View and Details Editors,” on page 25](#).

The Progress view can be toggled on and off by clicking the text located in the lower right of the VM Client window. If the Error Log view has also been opened, it toggles between the two.

You can double-click in the Progress view’s title area to expand it to use up the entire VM Client window. Double-click again to return it to its previous size and location.

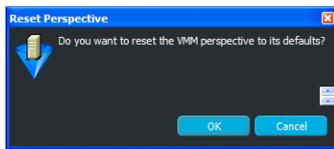
The  and  icons in the upper right of the Progress view are not used by the VM Client. Even though they might appear to work, their functions do nothing.

## 5.18 Resetting the VM Client Perspectives

You can reset the various perspectives to their original configuration:

- 1 In the **VM Client**, click *View > Reset Perspective*.

The following dialog box is displayed:



- 2 Click *OK* to reset the perspectives:
  - ♦ The Inventory panel is reset to its original size and location, to the left of the VM Client window.
  - ♦ The Details lists and editors are reset to their original location, to the right of the Inventory panel.
  - ♦ If the Error Log or Progress view is open, it is reset to the lower right portion of the window.

This action cannot be undone. You can reconfigure the perspectives again as needed.

## 5.19 Viewing and Editing VM Client Details

Review the following sections for information on the Details views:

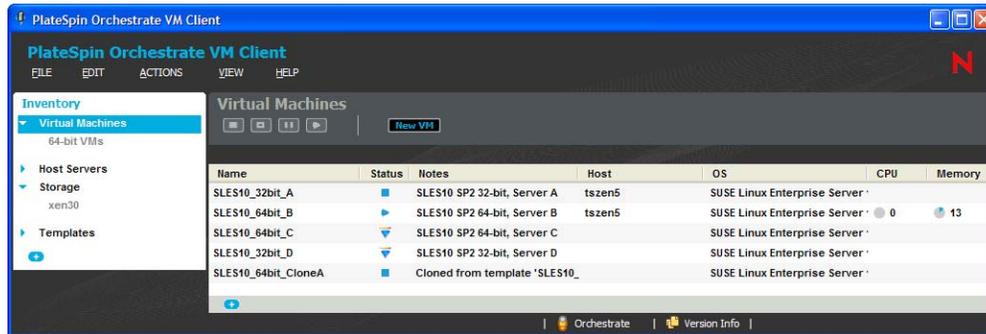
- ♦ [Section 5.19.1, “Virtual Machine Details,” on page 75](#)
- ♦ [Section 5.19.2, “Host Server Details,” on page 75](#)
- ♦ [Section 5.19.3, “Storage Details,” on page 76](#)
- ♦ [Section 5.19.4, “Template Details,” on page 76](#)
- ♦ [Section 5.19.5, “Details Tabs,” on page 77](#)

## 5.19.1 Virtual Machine Details

The details for each VM consist of a summary, monitoring information, and an event log.

- 1 In the **VM Client**, click the *Virtual Machines* view.

The VMs are displayed:



- 2 Do one of the following:

- ◆ Double-click one of the VMs.
- ◆ Right-click one of the VMs, then select *Show Details*.

The tabs for the selected VM are displayed.

- 3 For details on the tabs, see:

- ◆ “Summary Tab” on page 82
- ◆ “Monitoring Tab” on page 80
- ◆ “Event Log Tab” on page 79

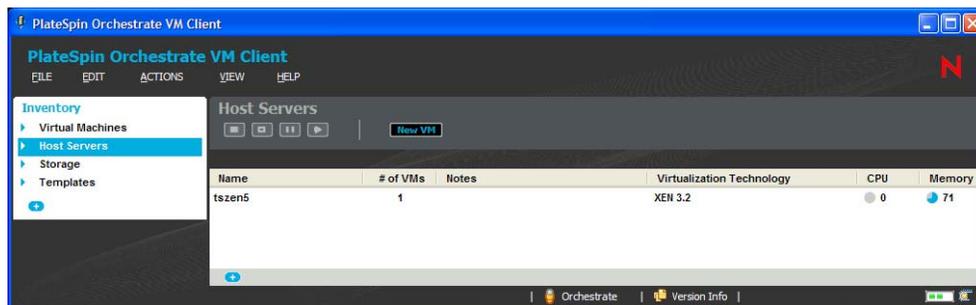
- 4 To close the details view, click the view name in the Inventory perspective panel.

## 5.19.2 Host Server Details

The details for each host server consist of a summary, configuration information, monitoring information, and an event log.

- 1 In the **VM Client**, click the *Host Servers* view.

The host servers are displayed:



- 2 Do one of the following:
  - ♦ Double-click one of the host servers.
  - ♦ Right-click one of the host servers, then select *Show Details*.

The tabs for the selected host are displayed.

- 3 For details on the tabs, see:
  - ♦ “Summary Tab” on page 82
  - ♦ “Configuration Tab” on page 77
  - ♦ “Monitoring Tab” on page 80
  - ♦ “Event Log Tab” on page 79

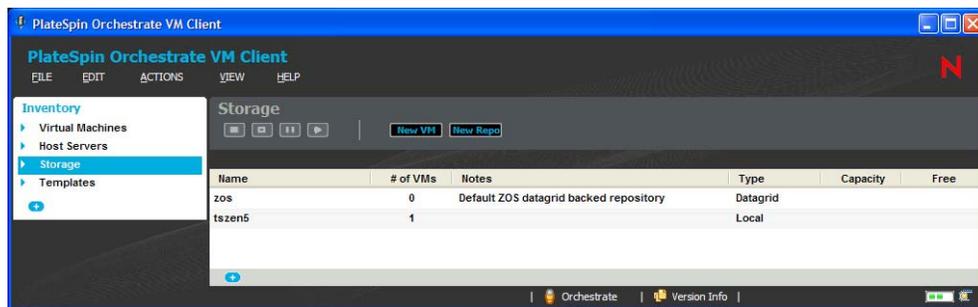
- 4 To close the details view, click the view name in the Inventory perspective panel.

### 5.19.3 Storage Details

The details for each storage item consist of a summary, VMs using this storage capacity, and host servers providing the storage space.

- 1 In the **VM Client**, click the *Storage* view.

The storage items are displayed:



- 2 Do one of the following:
  - ♦ Double-click one of the storage items.
  - ♦ Right-click one of the storage items, then select *Show Details*.

The tabs for the selected storage item are displayed.

- 3 For details on the tabs, see:
  - ♦ “Summary Tab” on page 82
  - ♦ “Virtual Machines Tab” on page 87
  - ♦ “Hosts Tab” on page 80

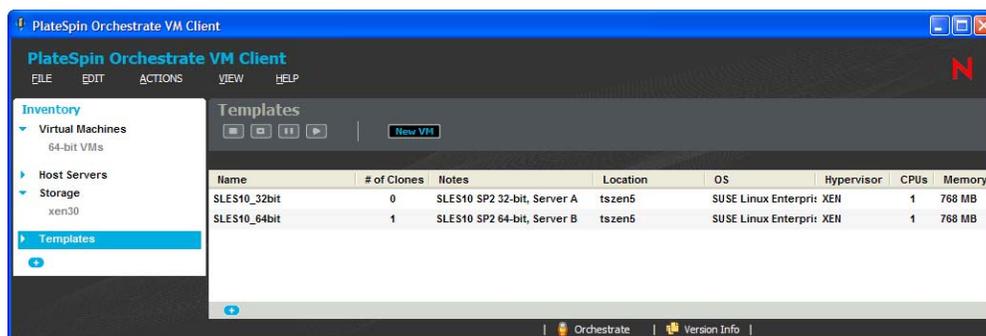
- 4 To close the details view, click the view name in the Inventory perspective panel.

### 5.19.4 Template Details

The details for each template consist of a summary and an event log.

- 1 In the **VM Client**, click the *Templates* view.

The templates are displayed:



2 Do one of the following:

- ◆ Double-click one of the templates.
- ◆ Right-click one of the templates, then select *Show Details*.

The tabs for the selected template are displayed.

3 For details on the tabs, see:

- ◆ “[Summary Tab](#)” on page 82
- ◆ “[Event Log Tab](#)” on page 79

4 To close the details view, click the view name in the Inventory perspective panel.

## 5.19.5 Details Tabs

The following sections provide information on the various details tabs:

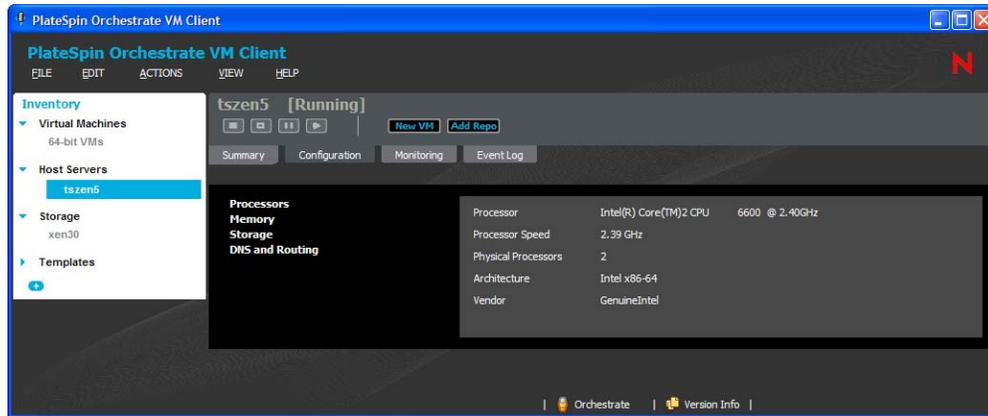
- ◆ “[Configuration Tab](#)” on page 77
- ◆ “[Event Log Tab](#)” on page 79
- ◆ “[Hosts Tab](#)” on page 80
- ◆ “[Monitoring Tab](#)” on page 80
- ◆ “[Summary Tab](#)” on page 82
- ◆ “[Virtual Machines Tab](#)” on page 87

### Configuration Tab

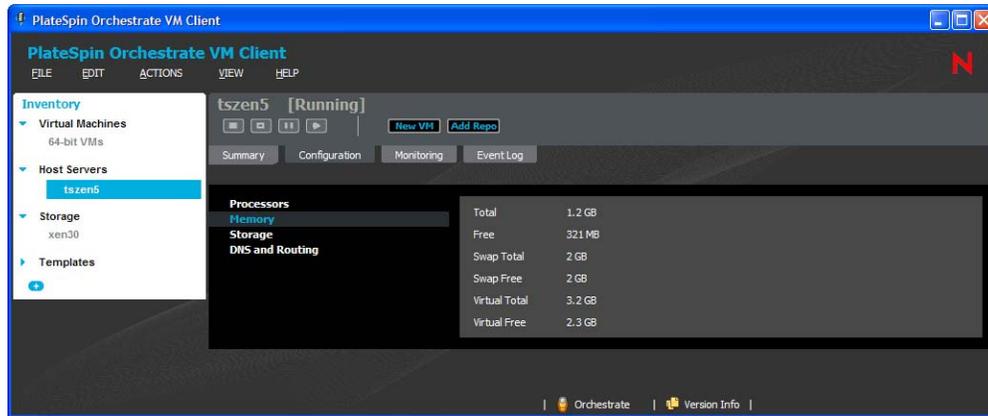
This details tab is used only by host servers. You use information available in the tab to effectively manage the existing VMs hosted on the VM hosts. For example, you can decide to stop VM when the memory utilization is high. You can also effectively utilize the VM host by deploying additional VMs, if required, and you can install applications on the VMs.

You can do the following with this tab:

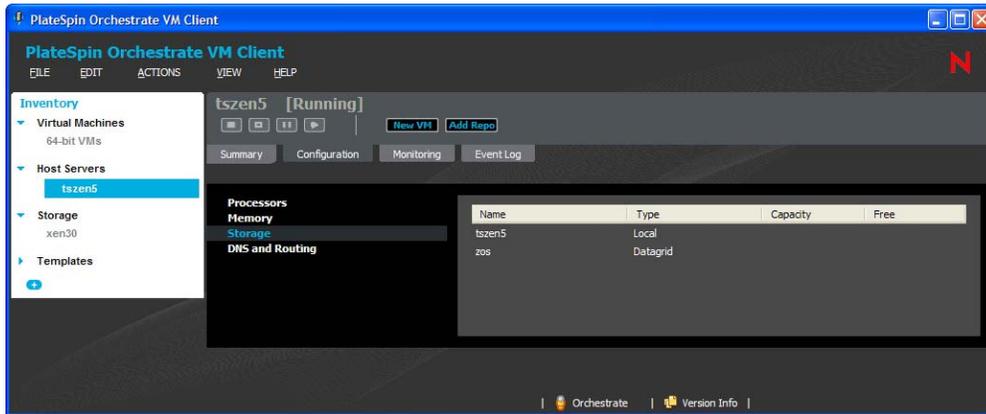
- ◆ Click *Processors* to view information about the processors on the host server. This is the default view:



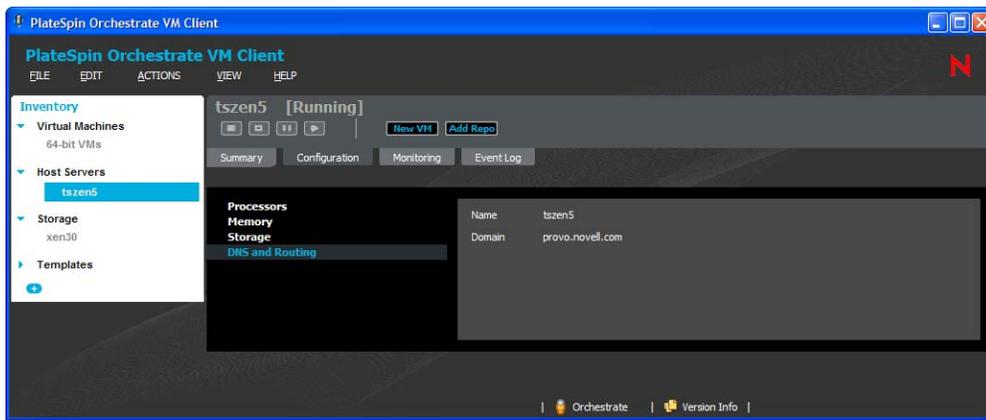
- ◆ Click *Memory* to view information about memory usage on the host server:



- ◆ Click *Storage* to view information about storage availability on the host server:



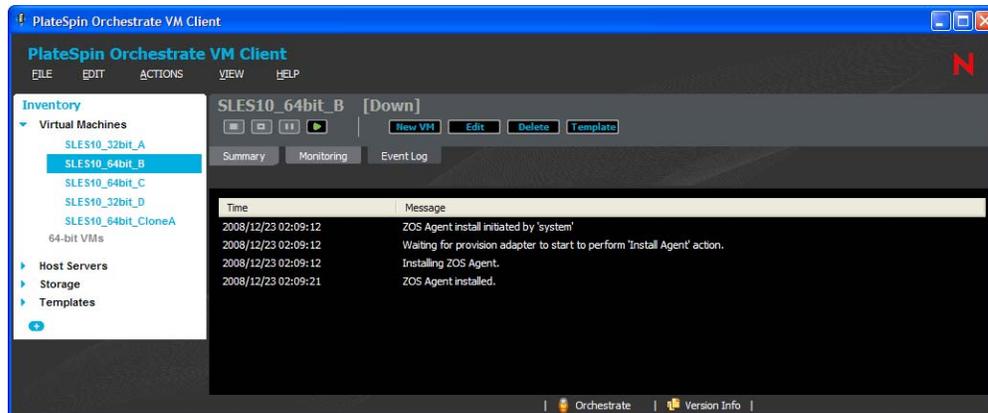
- ◆ Click *DNS and Routing* to view the computer name and domain name for the host server:



## Event Log Tab

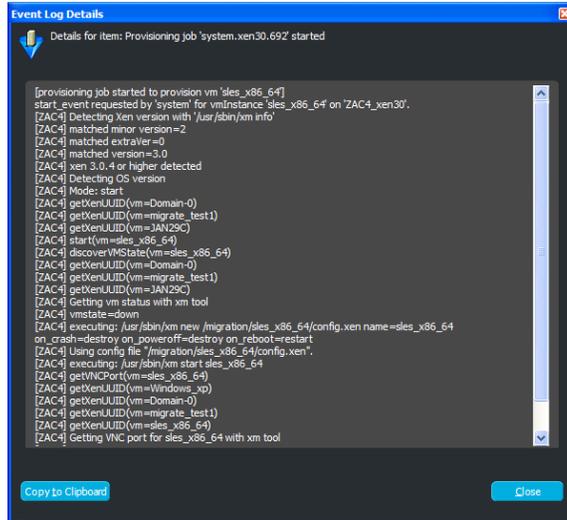
This details tab is used by VMs and host servers:

*Figure 5-2 Event Log Tab for VMs*



You can do the following with this tab:

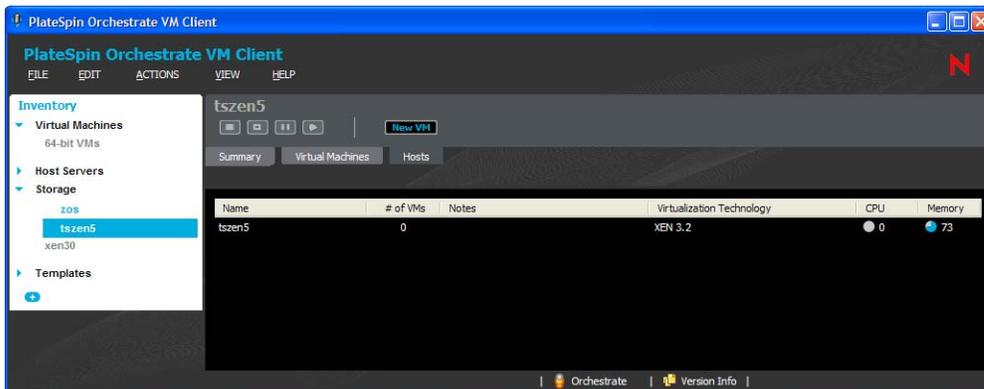
- ◆ View the Event Log information.
- ◆ Double-click an event item that contains the  icon to view the details of the event as reported by the Orchestrate Server. The following is an example of an event's details:



## Hosts Tab

This details tab is used only by storage items:

**Figure 5-3** Hosts Tab for Storage



You can do the following with this tab:

- ◆ View information about the host server this storage item is associated with.
- ◆ Double-click a listed host server to view its details tabs.

## Monitoring Tab

This details tab allows you to monitor many elements of both the virtual machines and physical host machines, aiding you in the overall management of your data center.

DNS must be working properly for this tab to display any information.

The Monitoring Agent must be installed on the device in order for this tab to show any information. For instructions to install the agent manually, see [Section 3.8, “Installing the Monitoring Agent on a VM,”](#) on page 39.

Review the following to understand the *Monitoring* tab and how to access and use it:

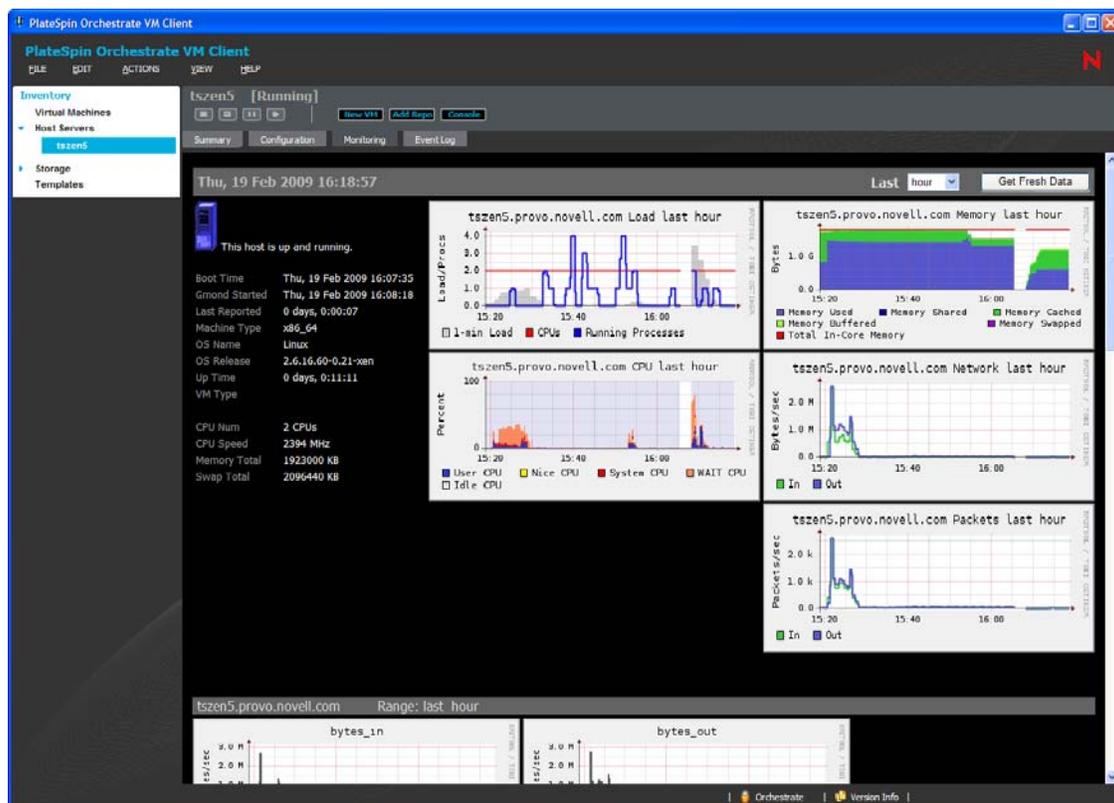
- ♦ “Understanding Monitoring” on page 81
- ♦ “Monitoring VMs and Host Servers” on page 81

## Understanding Monitoring

The *Monitoring* tab displays information about the VM or host server that you can use to monitor whether it is performing correctly, is overused or underused, its history, and details about its capacities.

The following is an example of the *Monitoring* tab’s information:

**Figure 5-4** Monitoring Tab for a Host Server



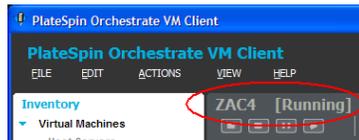
This example shows that information can be displayed side-by-side when you expand the width of the VM Client window.

## Monitoring VMs and Host Servers

- 1 In the **VM Client**, click the *Virtual Machines* view or the *Host Servers* view, then double-click a VM or host server that you want to monitor.

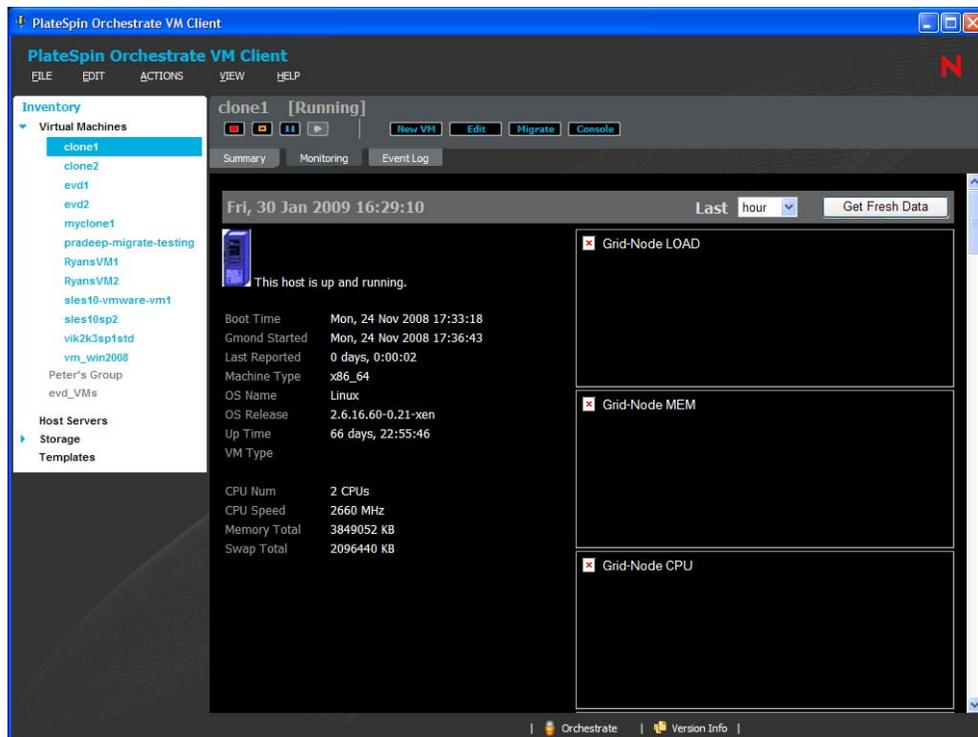
Wherever N/A appears in the details, it means that this information is not available.

- 2 Click the *Monitoring* tab to display the *Monitoring* view.
- 3 If you are requested to provide credentials, enter the Monitoring login name and password.
- 4 To set the time interval for which you want the data displayed (hour/day/week/month/year), go to the *Last* field and select the time from the drop-down options.
- 5 To maximize the *Monitoring* tab view in the VM Client window, double-click the the VM or Host name located just above the buttons:



Double-click again to return the tab to its previous view.

- 6 If the Monitoring view appears as follows, click the *Get Fresh Data* button:



You can refresh the view at any time.

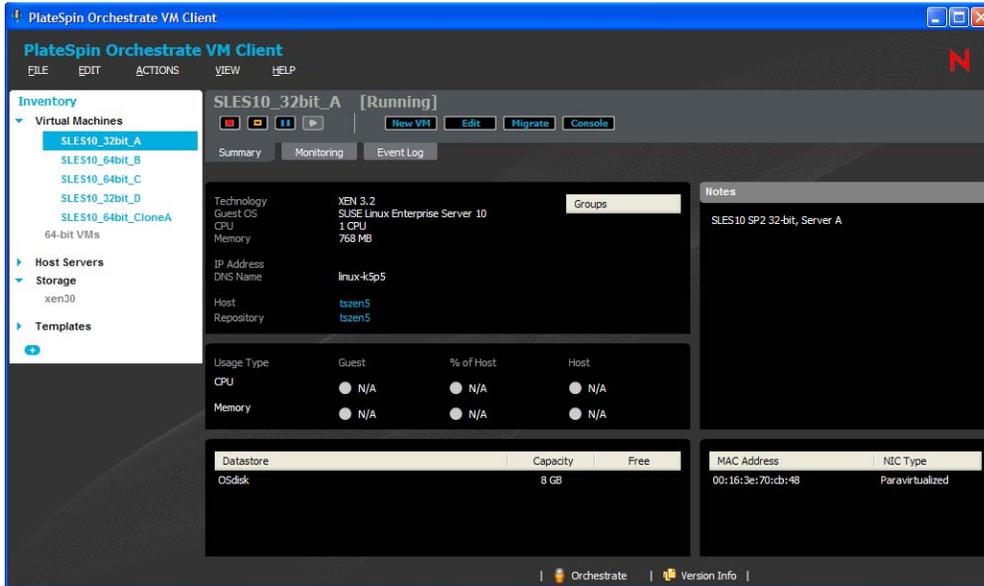
## Summary Tab

This details tab is used by each Inventory view type:

- ◆ “Virtual Machines Summary Tab” on page 83
- ◆ “Host Servers Summary Tab” on page 84
- ◆ “Storage Summary Tab” on page 85
- ◆ “Templates Summary Tab” on page 86

## Virtual Machines Summary Tab

Figure 5-5 Summary Tab for Virtual Machines



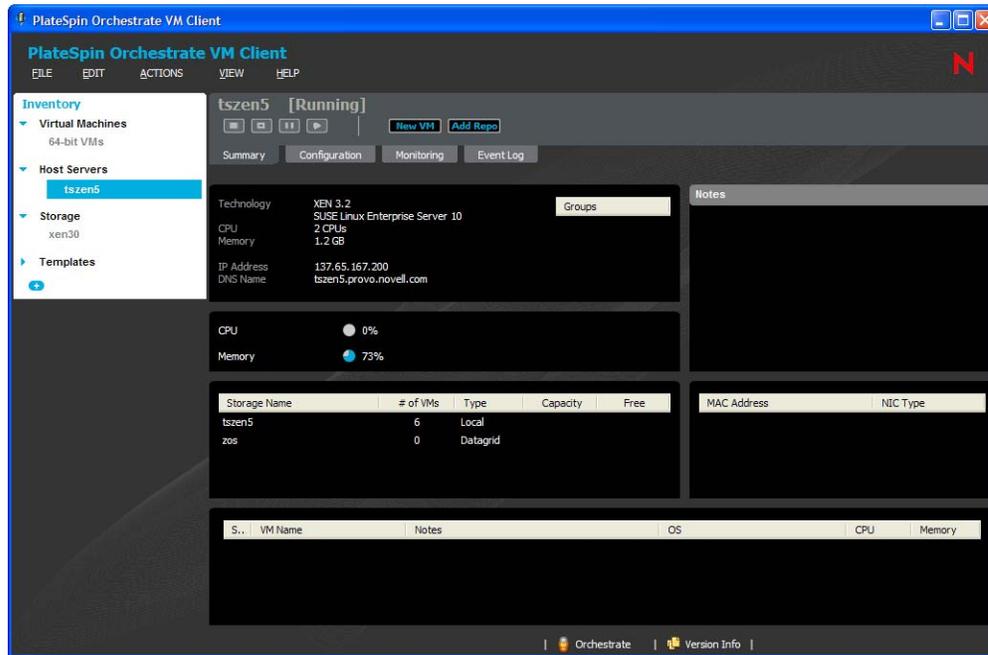
You can do the following with the *Summary* tab for VMs:

- ◆ In the first section, right-click a group name under *Groups*, then select *Remove* to revoke the VM's membership in the group.
- ◆ In the *Notes* section, enter and edit notes for the VM.
- ◆ View basic information about the VM, including:
  - ◆ The VM's operating system
  - ◆ The virtual machine technology being used by the VM
  - ◆ CPU and memory, including how it is used by the guest and host and the percentage of the host memory being used
  - ◆ IP address and DNS name, if it is not DHCP
  - ◆ Host and repository servers
  - ◆ Type of datastore and capacity
  - ◆ The VM's MAC address
- ◆ Double-click the *Host* entry to view that host server's details tabs.
- ◆ Double-click the *Repository* entry to view that storage item's details tabs.

Wherever N/A appears in the details, it means that the information is not available.

## Host Servers Summary Tab

Figure 5-6 Summary Tab for Host Servers



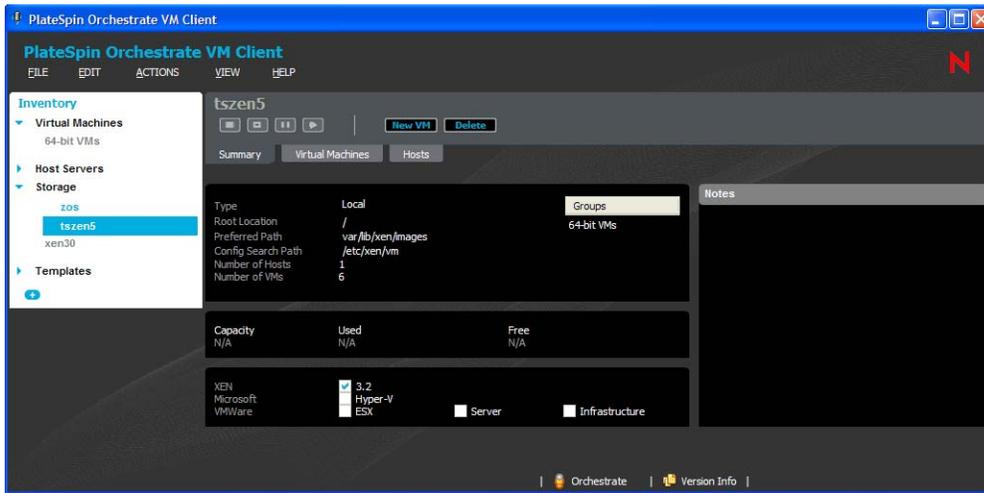
You can do the following with the *Summary* tab for host servers:

- ◆ In the first section, right-click a group name under *Groups*, then select *Remove* to revoke the host server's membership in the group.
- ◆ In the *Notes* section, enter and edit notes for the host server.
- ◆ View basic information about the host server, including:
  - ◆ The host server's operating system
  - ◆ The virtual machine technology being used by the host server
  - ◆ Percentage of CPU and memory in use
  - ◆ IP address and DNS name, if it is not DHCP
  - ◆ Storage items being used by the host server, including how many VMs are using each repository, the type of network access to the repository, and each repository's capacity and percent free
  - ◆ A list of VMs being hosted by this server
  - ◆ The host server's MAC address
- ◆ Right-click a listed storage item, then select either *Remove Repository* or *Associate Repository*.
- ◆ Double-click a listed storage item to view its details tabs.
- ◆ Right-click a listed VM, then select any of the available menu options. These are the same right-click menu options that are available for the VM in the *Virtual Machines* Inventory listing.
- ◆ Double-click a listed VM to view its details tabs.

Wherever N/A appears in the details, it means that the information is not available.

## Storage Summary Tab

Figure 5-7 Summary Tab for Storage Items



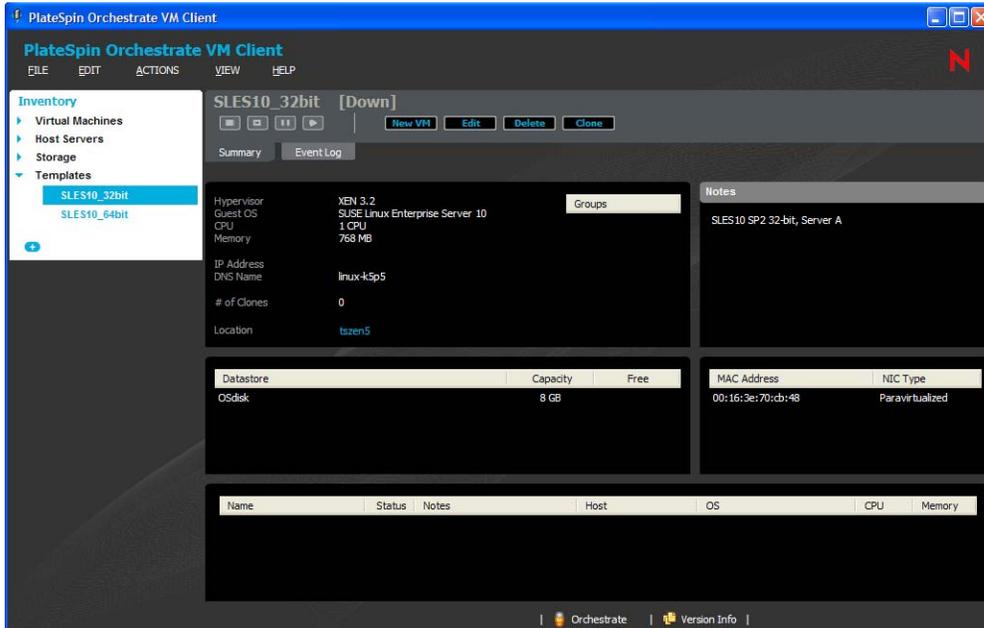
You can do the following with the *Summary* tab for storage items:

- ◆ In the first section, right-click a group name under Groups, then select *Remove* to revoke the storage item's membership in the group.
- ◆ In the Notes section, enter and edit notes for the storage item.
- ◆ View basic information about the storage item, including:
  - ◆ The repository's type and path information, including the ability to edit the information
  - ◆ The repository's capacity, including how much is being used and how much is free
  - ◆ Which hypervisors the location can be used with, including the ability to enable or disable them for this repository

Wherever N/A appears in the details, it means that the information is not available.

## Templates Summary Tab

Figure 5-8 Summary Tab for Templates



You can do the following with the *Summary* tab for templates:

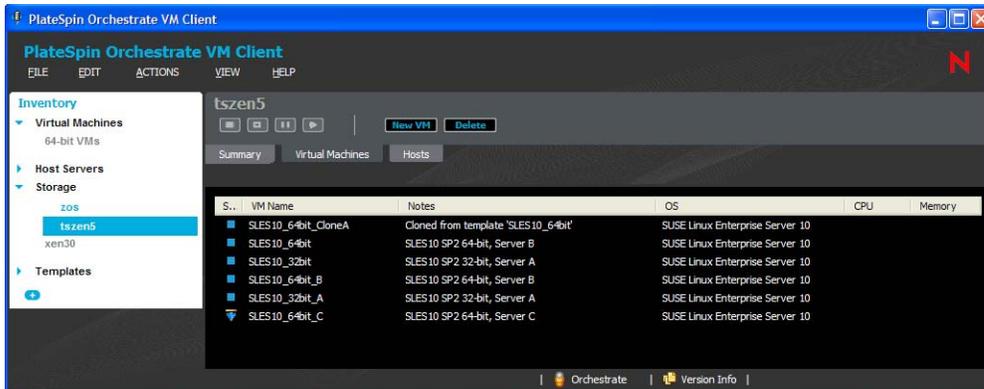
- ◆ In the first section, right-click a group name under *Groups*, then select *Remove* to revoke the template's membership in the group.
- ◆ In the *Notes* section, enter and edit notes for the template.
- ◆ View basic information about the template, including:
  - ◆ The template's operating system
  - ◆ The virtual machine technology being used by the template
  - ◆ Number of clones of the template
  - ◆ IP address and DNS name, if it is not DHCP
  - ◆ The type of datastore and capacity
  - ◆ The template's MAC address
  - ◆ The VMs it is associated with
- ◆ Double-click the *Location* entry to view the storage location's details tabs.
- ◆ Double-click a listed VM to view its details tabs. These are the VMs that were cloned from the template.
- ◆ Right-click a listed VM, then select any of the available menu options. These are the same right-click menu options that are available for the VM in the *Virtual Machines* Inventory listing.

Wherever N/A appears in the details, it means that the information is not available.

## Virtual Machines Tab

This details tab is used only by storage items:

**Figure 5-9** *Virtual Machines tab for Storage*



You can do the following with this tab:

- ◆ View which VMs are associated with the storage item.
- ◆ View the status of each associated VM.
- ◆ Double-click a listed VM to view its details tabs.



# Troubleshooting Virtual Machine Management

# 6

The following sections provide solutions to the problems you might encounter while working with the PlateSpin® Orchestrate VM Client:

- ♦ [Section 6.1, “Adding Multiple NPIV Disks to a VM,” on page 89](#)
- ♦ [Section 6.2, “Accessing VMs and Host Servers,” on page 89](#)
- ♦ [Section 6.3, “Cloning VMs,” on page 90](#)
- ♦ [Section 6.4, “Moving or Migrating the VMs,” on page 90](#)
- ♦ [Section 6.5, “Registering VM Hosts,” on page 91](#)

## 6.1 Adding Multiple NPIV Disks to a VM

- ♦ [“Duplicate NPIV disks are not allowed” on page 89](#)

### Duplicate NPIV disks are not allowed

Source: The PlateSpin Orchestrate VM Client.

Action: When you add an additional NPIV disk to a VM, ensure that the disk values for Fabric ID, World Wide Name, Target ID, and LUN are unique for each NPIV disk of the VM.

## 6.2 Accessing VMs and Host Servers

- ♦ [“VM or host server is not accessible” on page 89](#)

### VM or host server is not accessible

Source: The PlateSpin Orchestrate VM Client.

Explanation: A VM or host server does not appear to be accessible for any actions. This might be because its *Enabled* check box is disabled. By default, this check box is enabled.

Action: In the Development Client:

- 1 In the *Explorer* panel, expand the *Resources* section.
- 2 Browse for and select the VM or host server.
- 3 Click the *Info* tab and expand the *Info* section.
- 4 For a VM, expand its *Resource Information* section; for a host server, expand its *VmHost Information* section.
- 5 For a VM, make sure the *Resource Enabled* check box is selected; for a host server, make sure the *Enabled* check box is selected.
- 6 If you changed the status of the check box, click *File > Save* to make the change effective.

You should now be able to perform actions against the VM or host server.

## 6.3 Cloning VMs

- ♦ [“VMs Cloned on SLES Machines Might Not Appear on the Network” on page 90](#)

### VMs Cloned on SLES Machines Might Not Appear on the Network

Source: The PlateSpin Orchestrate VM Client.

Explanation: By default, SLES network configuration is set up with `FORCE_PERSISTENT_NAMES=yes` in `/etc/sysconfig/network/config`. This results in network device configurations being bound statically to specific MAC addresses.

MAC addresses in VM clones are dynamic. In particular, if you clone a VM, you must change the MAC address so that the new VM is unique on the local network segment.

If you have a SLES VM configured for DHCP with the default network configuration options, the clone does not appear on the network because its virtual NIC has a different MAC than the hard-coded configuration inside the VM image. The NIC looks like a new interface that isn't configured.

Action: Work around this issue by setting `FORCE_PERSISTENT_NAMES=no` in `/etc/sysconfig/network/config`.

This setting causes the networking configuration to revert to the traditional mode of assigning `eth0` to the first NIC detected by the kernel, `eth1` to the second, and so on. This is the preferred mode for VMs because a VM MAC address does not remain static.

In addition, a VM is likely to have only one or two virtual NICs, so `eth0`, `eth1`, and so on always refer to the same virtual NIC.

## 6.4 Moving or Migrating the VMs

- ♦ [“Moving or migrating VMs between two ESX hosts that are registered to a vCenter server by using the Orchestrate Development Client fails” on page 90](#)
- ♦ [“Moving a VM from one ESX host local storage to another ESX host local storage might fail” on page 91](#)
- ♦ [“A migrated VM's Orchestrate Agent loses connection with the Orchestrate Server” on page 91](#)

### Moving or migrating VMs between two ESX hosts that are registered to a vCenter server by using the Orchestrate Development Client fails

Source: The PlateSpin Orchestrate VM Client.

Action: Do the following:

- 1 Disconnect and remove one of the ESX hosts from the vCenter server.
- 2 Move or migrate the VMs by using the Orchestrate Development Client.

For more information, see [Section 4.2, “Moving VMs,” on page 53](#) or [Section 4.3, “Migrating VMs,” on page 54](#).

### **Moving a VM from one ESX host local storage to another ESX host local storage might fail**

Source: The PlateSpin Orchestrate VM Client.

Explanation: When you try to use the VM Client to move a VM of considerable size from one ESX host local storage to another ESX local storage, the move job might fail with the following error message:

```
Job timeout, because Max elapsed time expired.
```

Action: In the policy associated with the VM, appropriately increase the timeout value. For more information, see [“Provisioning a Virtual Machine” in the \*PlateSpin Orchestrate 2.0 Virtual Machine Management Guide\*](#).

### **A migrated VM’s Orchestrate Agent loses connection with the Orchestrate Server**

Source: The PlateSpin Orchestrate VM Client.

Explanation: After you migrate a Xen VM that has the Orchestrate Agent installed, the agent might lose connection with the Orchestrate Server.

Action: Restart the agent on the VM to reestablish the connection.

## **6.5 Registering VM Hosts**

- ♦ [“VM Hosts are not registered to the Orchestrate Server even after a considerable time” on page 91](#)

### **VM Hosts are not registered to the Orchestrate Server even after a considerable time**

Source: The PlateSpin Orchestrate VM Client.

Action: In the VM Client, click *Action > Discover Hosts*.

The registered VM host is displayed in the Host Servers listing.



# VM Installation Sources

# A

You have two sources from which to install VMs for the PlateSpin® Orchestrate VM Client:

- ♦ [Section A.1, “Disk Installation Sources,” on page 93](#)
- ♦ [Section A.2, “Guest Operating System Installation Sources,” on page 93](#)

## A.1 Disk Installation Sources

To install a fully virtualized machine, you can specify an installation source that must be mounted in such a way that the VM recognizes it as a mounted installation source. For example, FTP, HTTP, and NFS.

For other installation sources to use for both fully virtualized and paravirtualized installations, the following sources can be designated to serve as virtual mount points:

- ♦ A disk image, such as an ISO file accessible from the local file system of the machine performing the installation. For example: `/install_ISOs/sles10sp2.iso`

---

**IMPORTANT:** This ISO needs to be available on all Orchestrate Servers in the virtualization grid that were installed with the Virtual Machine Builder pattern, because a VM could access this ISO from any of them.

---

- ♦ A remote storage device specified by using the Internet SCSI (iSCSI) protocol. For example: `iscsi:iqn.2001-04.com.acme@0ac47ee2-216e-452a-a341-a12624cd0225`
- ♦ A remote storage device specified by using a Fibre Channel (NPIV) protocol. For example: `npiv:210400e08b80c40f`

Both Linux and Windows operating systems can use these installation sources, as applicable.

## A.2 Guest Operating System Installation Sources

You need to designate an installation source for a paravirtualized machine. The following is a list of guest operating systems and the installation sources to use for a paravirtualized installation. The listed sources are suggestions and examples.

- ♦ [Section A.2.1, “SUSE,” on page 93](#)
- ♦ [Section A.2.2, “NetWare,” on page 94](#)
- ♦ [Section A.2.3, “Other Linux Installation Sources,” on page 94](#)

### A.2.1 SUSE

- ♦ A remote directory containing the install files, specified by using the HTTP protocol: The HTTPS protocol is not supported.

`http://www.domain.com/install_sources/sles10sp2/`

- ♦ A remote directory containing the install files, specified by using the FTP protocol:

```
ftp://www.domain.com/install_sources/sles10sp2/
```

- ♦ A remote directory containing the install files, specified by using the NFS protocol:

```
nfs://www.domain.com/install_sources/sles10sp2/
```

## A.2.2 NetWare

A paravirtualized NetWare<sup>®</sup> operating system can be installed by using the same installation sources as fully virtualized VMs.

---

**NOTE:** Open Enterprise Server (OES) 2 and OES 2 Linux can be paravirtualized or fully virtualized. NetWare 6.5 and earlier must be fully virtualized. See [Section A.1, “Disk Installation Sources,” on page 93](#) for instructions on fully virtualizing the guest operating system.

---

## A.2.3 Other Linux Installation Sources

Paravirtualized operating systems other than SUSE<sup>®</sup> and NetWare can be installed by using the same installation sources as the fully virtualized VMs (see [Section A.1, “Disk Installation Sources,” on page 93](#)).

# NPIV Terminology

# B

Terms associated with N\_Port ID Virtualized (NPIV) Storage:

- ♦ **Fabric ID:** A unique numeric identifier of the SAN repository.
- ♦ **LUN:** Logic Unit Number. A unique number used to identify a logical unit on the disk storage. The value can be between 0 and 254.
- ♦ **Target ID:** The port number of the storage disk provided by the SAN vendor.
- ♦ **WWN:** World Wide Name. An 8-byte number used as a unique identifier in a Fibre Channel or Serial Attached SCSI storage network. The first 16 digits represent the WWPN; the last 16 digits represent the WWNN.
- ♦ **WWNN:** World Wide Node Name. A 16-digit World Wide Name assigned to a node (an endpoint or a device) in a Fibre Channel fabric or Serial Attached SCSI storage network.
- ♦ **WWPN:** World Wide Port Name. A 16-digit World Wide Name assigned to a port in a Fibre Channel fabric or Serial Attached SCSI storage network.



# Index

# C

This short index is designed to provide quick access to hard-to-find information in the *PlateSpin® Orchestrate VM Client Reference and Guide*. Some index entries are to other PlateSpin Orchestrate guides for information related to managing VMs.

An extensive index is not intended here, so most of the section titles that are visible on the Contents page for this manual are not included.

- Agent, Monitoring
- Agent, Orchestrate
- Canceling a VM installation
- Configuration Details dialog box
- Configuration details tab
- Development Client
- Discovered (#) group for host servers
- Enabled check box in the Development Client
- Event Log details tab
- Groups
- Hosts details tab
- Installing the VM Client
- Migrating VMs
- Moving VMs
- Monitoring Agent
- Monitoring details tab
- Orchestrate Agent
- Provisioning adapter policies
- Provisioning in the Development Client
- Repositories, adding
- Repositories, associating
- Repositories, managing in the Development Client
- Repositories, removing
- Reset Perspective menu option
- Saving changes
- Sorting
- Summary details tab
- Version Information dialog box
- Virtual Machines details tab
- Virtualization host types
- VM technologies and actions
- VNC
- Windows VM installation source

We encourage you to use the User Comment feature on the [PlateSpin Orchestrate documentation Web site \(http://www.novell.com/documentation/pso\\_orchestrate20/\)](http://www.novell.com/documentation/pso_orchestrate20/) to help us keep this index up-to-date and useful for you.

To use the feature, simply access [this Index section \(http://www.novell.com/documentation/pso\\_orchestrate20/pso20\\_vmclient/data/bj2pd1o.html\)](http://www.novell.com/documentation/pso_orchestrate20/pso20_vmclient/data/bj2pd1o.html) on the documentation Web site and suggest new entries or corrections to existing entries in the User Comment field there.

Please include the following information:

- ◆ The new entry's label and the URL to the section containing the related information.  
  
The proper URL can be obtained by right-clicking on the Web page link to that section, saving it to your operating system's clipboard, then pasting it into the User Comment.
- ◆ The existing entry's label and the correction for it.