

Development Client Reference

Novell® PlateSpin® Orchestrate

2.5

November 29, 2010

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 © 2008-2010 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
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

About This Guide	7
1 Layout	9
2 Orchestrate Development Client Menus and Tools	13
2.1 The Operations Menu Bar	13
2.1.1 File	13
2.1.2 Edit	14
2.1.3 View	18
2.1.4 Actions	18
2.1.5 Provision	19
2.1.6 Server	20
2.1.7 Windows	24
2.1.8 Help	24
2.2 The Orchestrate Development Client Toolbar	24
3 The Orchestrate Server and the Server Admin Objects	27
3.1 The Orchestrate Server Object	27
3.1.1 The Orchestrate Server Info/Configuration Page	28
3.1.2 The Orchestrate Server Authentication Page	34
3.1.3 The Orchestrate Server Policies Page	38
3.1.4 The Orchestrate Server Constraints/Facts Page	38
3.2 The Server Admin Object	39
4 The Job Object	41
4.1 Job Groups	41
4.2 The Job Info/Groups Tab	41
4.2.1 Info	42
4.2.2 Groups	51
4.3 The JDL Editor Tab	51
4.4 The Job Library Editor Tab	52
4.5 The Job Policies Tab	53
4.6 The Job Constraints/Facts Tab	54
5 The Resource Object	55
5.1 Resource Groups	55
5.2 The Resource Info/Groups Page	55
5.2.1 The Info Panel	55
5.2.2 The Groups Panel	82
5.3 The Provision Info Page	82
5.4 The Resource Log Page	83
5.5 The Resource Policies Page	83
5.6 The Resource Health Debugger Page	83
5.7 The Resource Constraints/Facts Page	84
5.8 Resource Object Naming and Renaming	84

6	The VM Host Object	85
6.1	The Info Page	85
6.1.1	Show Inherited Fact Values Check Box	86
6.1.2	VM Host Information Panel	86
6.1.3	Provisioning Adapter Config Panel	88
6.1.4	Guest VM Monitor Information Panel.	89
6.2	The Policies Page	90
6.3	The Health Debugger Page	90
6.4	The Constraints/Facts Page.	90
6.5	The Action History Page	90
6.6	VM Host Object Naming and Renaming	91
6.7	Unique VM Host Cluster Facts	91
6.7.1	Orchestrate Facts in the VM Host Cluster Object	91
6.7.2	Orchestrate Facts in a VM Host Residing in a Cluster	92
6.7.3	Orchestrate Facts in VMs Hosted in Clusters	93
6.8	vCPU Slots for VM Hosts	94
6.8.1	Configuring vCPUs on VM Hosts.	94
6.8.2	Configuring vCPUs on VM Host Clusters	95
6.8.3	Configuring vCPUs on VMs.	95
7	The Virtual Disk Object	97
7.1	Understanding the Virtual Disk Object	97
7.1.1	Creating Or Deleting a vDisk in the Development Client	97
7.1.2	Sharing Virtual Disks Among VM Hosts	100
7.1.3	Moving Virtual Disks	100
7.2	Viewing Virtual Disk Configuration in the Development Client.	102
7.2.1	The Virtual Disk Information Panel	102
7.2.2	The Virtual Disk Policies Tab	104
7.2.3	The Virtual Disk Health Debugger Tab	104
7.2.4	The Virtual Disk Constraints/Facts Tab	105
7.2.5	Virtual Disk Object Naming and Renaming	105
8	The Virtual NIC Object	107
8.1	Understanding the Virtual NIC Object	107
8.1.1	The Purpose of the Virtual NIC	107
8.1.2	Creating Or Deleting a vNIC in the Development Client	107
8.2	Viewing the Virtual NIC Configuration in the Development Client	110
8.2.1	The Virtual NIC Info Panel.	111
8.2.2	The Virtual NIC Policies Tab	115
8.2.3	The Virtual NIC Health Debugger Tab.	115
8.2.4	The Virtual NIC Constraints/Facts Tab	115
8.2.5	Virtual NIC Object Naming and Renaming	115
9	The Network Group and its Virtual Bridge Objects	117
9.1	Understanding the Network Group and Virtual Bridge Objects	117
9.1.1	The Virtual Bridge Object.	117
9.1.2	The Purpose of the Virtual Bridge	118
9.1.3	Creating or Deleting a vBridge in the Development Client	118
9.1.4	Virtual Bridge Object Naming and Renaming	120
9.2	Viewing the Virtual Bridge Configuration in the Development Client	121
9.2.1	The Virtual Bridge Info/Groups Tab	122
9.2.2	The Virtual Bridge Policies Tab	123

9.2.3	The Virtual Bridge Health Debugger Tab	123
9.2.4	The Virtual Bridge Constraints/Facts Tab	123
10	The Repository Object	125
10.1	Right-Click Menu Actions on the Repository Object	125
10.2	Repository Groups	126
10.3	The Repository Info/Groups Tab	126
10.3.1	The Info Panel	126
10.3.2	Best Practices for Entering Repository File Paths	131
10.3.3	Groups	132
10.4	The Repository Policies Tab	132
10.5	The Repository Health Debugger Tab	133
10.6	The Repository Constraints/Facts Tab	133
10.7	The Repository Action History Tab	133
10.8	Repository Object Naming and Renaming	133
11	The User Object	135
11.1	User Groups	135
11.2	The User Info/Groups Tab	135
11.2.1	Info	135
11.2.2	Groups	141
11.3	The User Policies Tab	141
11.4	The User Health Debugger Tab	141
11.5	The User Constraints/Facts Tab	142
11.6	The User Action History Tab	142
12	Miscellaneous Objects Displayed in the Explorer Tree	143
12.1	The Policy Object	143
12.1.1	Policy Constraints	143
12.1.2	Policy Facts	143
12.2	Computed Fact Objects	143
12.3	Event Objects	144
12.4	Metrics Objects	144
13	The PlateSpin Orchestrate Job Scheduler	145
13.1	Understanding the Job Scheduler View	145
13.1.1	Navigating The Job Schedules Table	146
13.1.2	Creating or Modifying a Job Schedule	148
13.1.3	Understanding Cron Syntax in the Job Scheduler	158
13.2	Walkthrough: Scheduling a System Job	162
13.2.1	Deploying a Sample System Job	162
13.2.2	Creating a New Schedule for the Job	164
13.2.3	Defining the New Schedule	165
13.2.4	Activating the New Schedule	172
13.2.5	Running the New Schedule Immediately	172
14	The Policy Debugger	175
14.1	The Constraints Table View	175
14.1.1	The Match Context Area	176
14.1.2	The Constraint Type List	178

14.1.3	The Verbose Check Box	178
14.1.4	The Capable Resources Summary	178
14.1.5	The Constraints Column of the Constraints Table View	178
14.1.6	The Policy Column of the Constraints Table	180
14.2	The Facts Table View	180
14.2.1	The All Facts Check Box	181
14.3	Policy Debugger Use Cases	182
14.3.1	Use Case 1: Determining Why a Job is in a Waiting State	182
A	Grid Object Health Monitoring	185
A.1	Health Facts	185
A.2	Health Events	187
A.3	The Health Debugger	187
A.3.1	The Constraints Table Panel	188
A.3.2	The Facts Table View	191
B	Events	193
B.1	Event Object Visualization and Management in the Development Client	193
B.1.1	Deploying a New Rule-Based Event	194
B.1.2	Deploying a Pre-written Rule-Based Event	194
B.1.3	Undeploying an Event	195
B.1.4	The Event Editor	195
B.2	The Event Debugger	196
B.2.1	The Constraints Table	197
B.2.2	The Facts Table	199
B.3	Understanding the PlateSpin Orchestrate Events System	200
B.3.1	Event Notification	201
B.3.2	Built-in Events	201
B.3.3	Rule-based Events	202
C	The Metrics Facility	205
C.1	Metrics Facility Functionality	205
C.2	Ganglia Metrics	205
C.3	How Does the Metrics Facility Impact Orchestrate Server Performance?	207
C.3.1	I/O Contention	207
C.3.2	Too Many Open Files	207
C.4	RRD Definition Using Deployable .metric Files	207
C.4.1	XML Format for Deployable .metric Definitions	208
C.5	Query of Aggregated Metric Values	209
C.5.1	Example of a JDL Query for Aggregated Metric Values	209
C.5.2	Example of a Policy Constraint or Event Constraint Using Aggregated Metric Values	209
C.5.3	Example of Using Non-aggregated ("Raw") Historical Metric Values	210
C.6	MetricsManager MBean API	210
C.6.1	MBean Methods Exposed by the MetricsManager Facility	210
C.6.2	The MetricsDeployer Facility	211
C.7	Using the Metrics Facility in the Development Client Interface	211
D	Documentation Updates	213
D.1	November 29, 2010	213

About This Guide

This *PlateSpin Orchestrate Development Client Reference* introduces the PlateSpin Orchestrate Development Client, the product's basic administration environment. The guide provides an introductory overview of the Development Client interface. The guide is organized as follows:

- ♦ Chapter 1, "Layout," on page 9
- ♦ Chapter 2, "Orchestrate Development Client Menus and Tools," on page 13
- ♦ Chapter 3, "The Orchestrate Server and the Server Admin Objects," on page 27
- ♦ Chapter 4, "The Job Object," on page 41
- ♦ Chapter 5, "The Resource Object," on page 55
- ♦ Chapter 6, "The VM Host Object," on page 85
- ♦ Chapter 7, "The Virtual Disk Object," on page 97
- ♦ Chapter 8, "The Virtual NIC Object," on page 107
- ♦ Chapter 9, "The Network Group and its Virtual Bridge Objects," on page 117
- ♦ Chapter 10, "The Repository Object," on page 125
- ♦ Chapter 11, "The User Object," on page 135
- ♦ Chapter 12, "Miscellaneous Objects Displayed in the Explorer Tree," on page 143
- ♦ Chapter 13, "The PlateSpin Orchestrate Job Scheduler," on page 145
- ♦ Chapter 14, "The Policy Debugger," on page 175
- ♦ Appendix A, "Grid Object Health Monitoring," on page 185
- ♦ Appendix B, "Events," on page 193
- ♦ Appendix C, "The Metrics Facility," on page 205
- ♦ Appendix D, "Documentation Updates," on page 213

Audience

This book is intended for data center managers and IT or Operations administrators. It assumes that users of the product have the following background:

- ♦ General understanding of network operating environments and systems architecture.
- ♦ Knowledge of basic UNIX shell commands 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>) and enter your comments there.

Additional Documentation

In addition to this *PlateSpin Orchestrate Development Client Reference*, PlateSpin Orchestrate 2.5 documentation includes the following additional guides that contain valuable information about the product:

- ♦ *PlateSpin Orchestrate 2.5 Getting Started Reference*
- ♦ *PlateSpin Orchestrate 2.5 Installation and Configuration Guide*
- ♦ *PlateSpin Orchestrate 2.5 Upgrade Guide*
- ♦ *PlateSpin Orchestrate 2.5 VM Client Guide and Reference*
- ♦ *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*
- ♦ *PlateSpin Orchestrate 2.5 Administrator Reference*
- ♦ *PlateSpin Orchestrate 2.5 High Availability Configuration Guide*
- ♦ *PlateSpin Orchestrate 2.5 Command Line Reference*
- ♦ *PlateSpin Orchestrate 2.5 Server Portal Reference*
- ♦ *PlateSpin Orchestrate 2.5 Developer Guide and Reference*

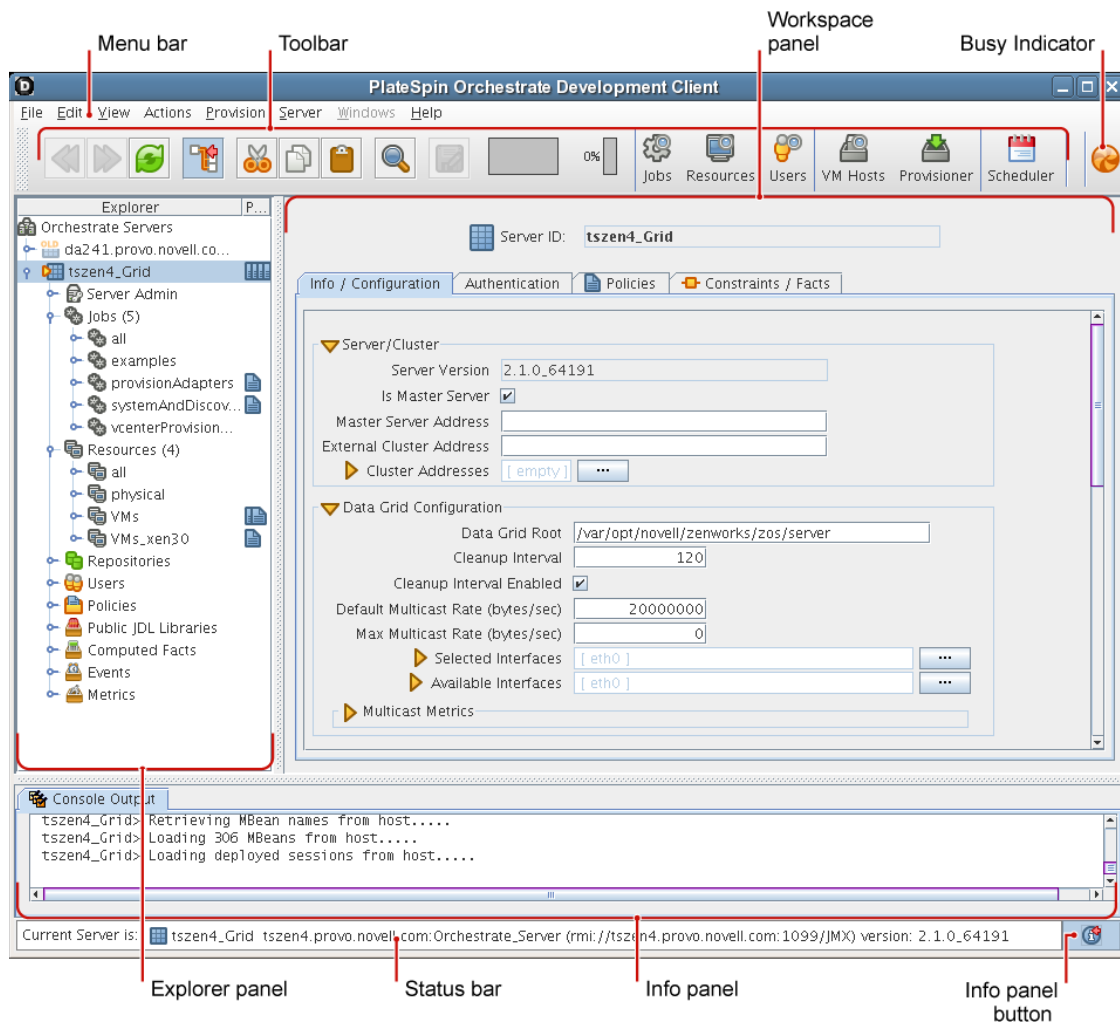
Layout

1

Both the grid administrator and the job developer need to have access to and use the PlateSpin Orchestrate Development Client. The administrator needs to use the console to perform any management functions, such as creating user accounts and managing Orchestrator Server activities. The developer uses the console to access the JDL editor for creating or modifying jobs and policies.

The following figure shows the general areas on the console interface that are referred to in this guide.

Figure 1-1 The PlateSpin Orchestrate Development Client with Parts Identified



The following chart describes the functional areas of the main PlateSpin Orchestrate Development Client display.

Table 1-1 Detailed Description of Console Areas

Area	Description
Menu bar	<p>Provides operations categorized under menus such as <i>File</i>, <i>Edit</i>, <i>View</i>, <i>Grid</i>, <i>Server</i>, <i>Windows</i>, and <i>Help</i>.</p> <ul style="list-style-type: none"> ♦ The <i>File</i> menu lets you save any changes you've made or exit the console. ♦ The <i>Edit</i> menu lets cut, copy, and paste items and choose general and server preferences for console. ♦ The <i>View</i> menu lets you manipulate the display of the different components of the console and refresh the Explorer and Workspace panels. ♦ The <i>Actions</i> menu lets you launch specific tools that create and delete users or user groups, computing resources, jobs, policies, and computed facts. ♦ The <i>Server</i> menu lets you start a local server, log in to the server, create and display logs for logged in servers, log out from the server, and shut down the server. ♦ The <i>Windows</i> menu lets you select console windows to display when you have more than one console window open. You can open the Explorer panel and the two tabs of the Info panel (<Orchestrator> <i>Log</i> and <i>Console Output</i>) in their own windows by right-clicking the tab and choosing <i>Open in window</i> in the pop-up menu. ♦ The <i>Help</i> menu provides access to the About box for the console. It also provides a link to ZENworks Orchestrator documentation on the Web.
Main toolbar	<p>The main toolbar has buttons for executing common tasks. The basic tasks are <i>Go Back</i>, <i>Go Forward</i>, <i>Refresh the view</i>, <i>Hide or Show the Explorer Panel</i>, <i>Cut</i>, <i>Copy</i>, <i>Paste</i>, and <i>Save changes in workspace view</i> and <i>Open the Find Dialog</i>.</p> <p>The toolbar also includes buttons that open monitoring views for <i>Jobs</i>, <i>Resources</i>, and <i>Users</i>.</p> <p>To the far left of the toolbar, a pinwheel icon indicates when the console is busy.</p>
Explorer panel	<p>The Explorer panel displays a hierarchical tree. The tree lets you navigate to different objects; you can click items in the tree to see their details. For example, you can display computing resources for a selected grid. When you click <i>Computing Resources</i> in the tree, its details appear in the Workspace panel with a list of active computing resources. You can edit the <i>Computing Resource</i> attributes in the workspace panel.</p>
Workspace panel	<p>The Workspace panel displays a detailed view for an item you select in the Explorer panel. For example, if you select a computing resource under <i>physical</i> in the Explorer panel, the Workspace panel view changes to show the details for that resource. You can edit the properties of an Orchestrator object in the views displayed in the Workspace panel.</p>
Info panel	<p>The Info panel displays a variety of information, such as validation and error messages, log files, and query results. You can display or hide the Info panel by clicking the <i>Info panel</i> button in the Status bar.</p>

Area	Description
Status bar	The status bar displays general identity information about the Orchestrator Server where you are logged in.

For information about launching the client and using it for the first time, see “[Walkthrough: Launching the PlateSpin Orchestrate Development Client](#)” in the *PlateSpin Orchestrate 2.5 Installation and Configuration Guide*.

For information about the user interface of the PlateSpin Orchestrate Development Client, see [Chapter 2, “Orchestrate Development Client Menus and Tools,”](#) on page 13.

Orchestrate Development Client Menus and Tools

2

A number of operations are available from the PlateSpin Orchestrate Development Client and can be accessed from its menu bar and toolbar.

- ♦ [Section 2.1, “The Operations Menu Bar,” on page 13](#)
- ♦ [Section 2.2, “The Orchestrate Development Client Toolbar,” on page 24](#)

2.1 The Operations Menu Bar

The Operations Menu Bar in the Orchestrate Development Client provides options that help you to create and administer objects in the Explorer Tree.

- ♦ [Section 2.1.1, “File,” on page 13](#)
- ♦ [Section 2.1.2, “Edit,” on page 14](#)
- ♦ [Section 2.1.3, “View,” on page 18](#)
- ♦ [Section 2.1.4, “Actions,” on page 18](#)
- ♦ [Section 2.1.5, “Provision,” on page 19](#)
- ♦ [Section 2.1.6, “Server,” on page 20](#)
- ♦ [Section 2.1.7, “Windows,” on page 24](#)
- ♦ [Section 2.1.8, “Help,” on page 24](#)

2.1.1 File

The File menu (Alt+F) provides keyboard and mouse accessible methods for users to save changes or to exit the application.

- ♦ [“Save” on page 13](#)
- ♦ [“Exit” on page 13](#)

Save

The Save operation provides a mouse and keyboard (File > Ctrl+S) accessible method for users to save any changes made in the visible view.

Exit

The exit operation provides a mouse and keyboard (File > Alt+X) accessible method for users to close all server connections and to exit the Orchestrate Development Client application.

2.1.2 Edit

The Edit menu (Alt+E) provides keyboard and mouse accessible methods for users to save changes or to exit the application.

- ♦ [“Undo Addition” on page 14](#)
- ♦ [“Redo” on page 14](#)
- ♦ [“Cut” on page 14](#)
- ♦ [“Copy” on page 14](#)
- ♦ [“Paste” on page 14](#)
- ♦ [“Find” on page 15](#)
- ♦ [“Find Next” on page 15](#)
- ♦ [“Find Previous” on page 15](#)
- ♦ [“Enter Find String” on page 15](#)
- ♦ [“Load Text” on page 15](#)
- ♦ [“Save Text” on page 15](#)
- ♦ [“Preferences” on page 16](#)

Undo Addition

The Undo operation provides a mouse-accessible method for users to undo the action they have just performed in the Orchestrate Development Client. The operation can also be executed from the keyboard (Ctrl+Z).

Redo

The Redo operation provides a mouse-accessible method for users to redo the action they have just performed in the Orchestrate Development Client. The operation can also be executed from the keyboard (Ctrl+Y).

Cut

The Cut operation provides a mouse-accessible method for users to cut the selected object and move it to the clipboard. The operation can also be executed from the keyboard (Ctrl+X).

Copy

The Copy operation provides a mouse-accessible method for users to copy the selected object to the clipboard. The operation can also be executed from the keyboard (Ctrl+C).

Paste

The Paste operation provides a mouse-accessible method for users to paste the contents of the clipboard to the desired location. The operation can also be executed from the keyboard (Ctrl+V).

Find

The Exit operation provides a mouse-accessible method for users to open the Find and Replace dialog box, where they can search for and replace (if necessary) editable strings located in logs and editing views (for example, the Policy Editor).

Figure 2-1 The Find and Replace Dialog Box Invoked From the Policy Editor



The operation can also be executed from the keyboard (Ctrl+F).

Find Next

The Find Next operation provides a mouse-accessible method for users to find the next occurrence of the string they previously searched for. The operation can also be executed from the keyboard (F3).

Find Previous

The Find Previous operation provides a mouse-accessible method for users to find the previous occurrence of the string they searched for. The operation can also be executed from the keyboard (Shift+F3).

Enter Find String

The Enter Find String operation provides a mouse-accessible method for users to load the text of the string they want to search for. The operation can also be executed from the keyboard (Ctrl+E).

Load Text

The Load Text operation provides a method for users to load text from an existing file into the open, editable view. When selected, the operation opens a browse dialog box where the file can be selected.

Save Text

The Save Text operation provides a method for users to save text in an editable, active view to a file. When selected, the operation opens a save dialog box where you can browse to a network location where you want to save the file. By default, the file is named according to the view and the context within which you are viewing it. You can change the name of the file when you save it.

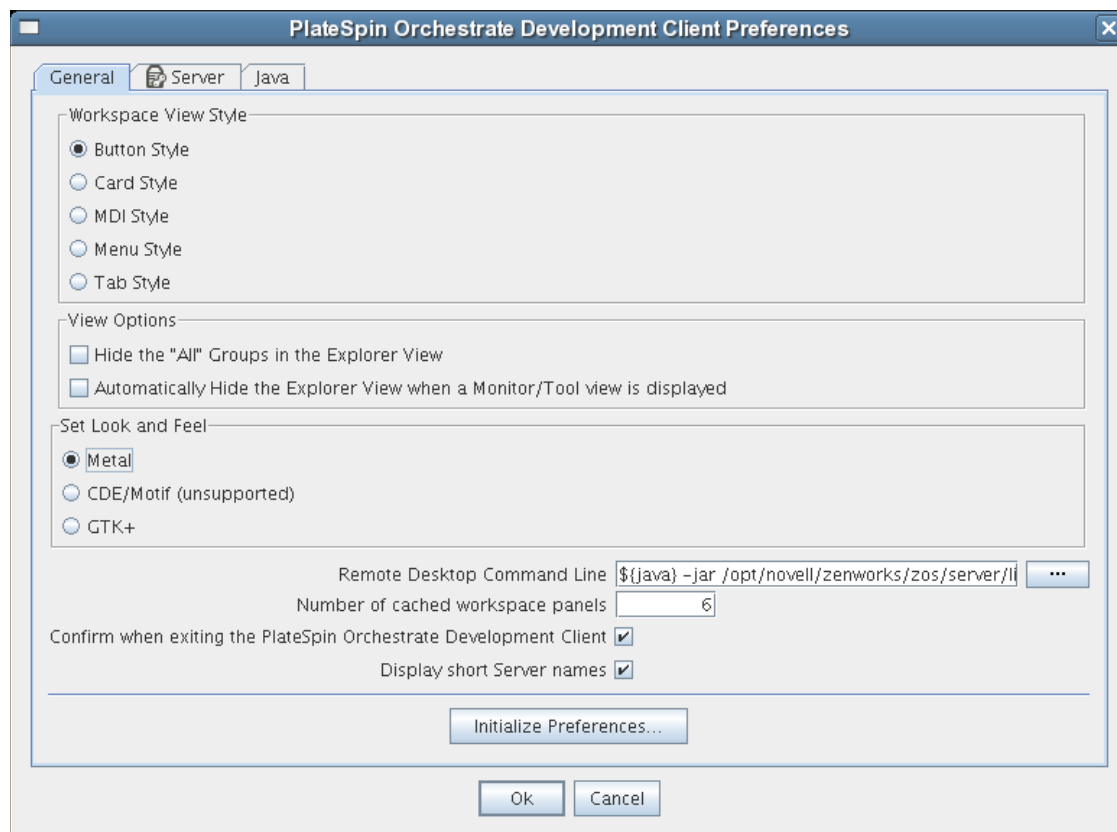
Preferences

The Preferences operation provides a method for users to change the preferences for the Orchestra Development Client display. When selected, the operation opens the Orchestra Development Client Preferences dialog box.

The dialog box has three tabbed pages.

General Page

Figure 2-2 General Page of the Orchestra Development Client Preferences



Preference settings on this page that you can change are self-explanatory. If you click *Initialize Preferences*, the preference settings (except *Look and Feel* settings) are initialized to installation values.

Server Page

Figure 2-3 Server Page of the Orchestrate Development Client Preferences

The screenshot shows the 'PlateSpin Orchestrate Development Client Preferences' dialog box with the 'Server' tab selected. The dialog has three tabs: 'General', 'Server', and 'Java'. The 'Server' tab contains four sections: 'Log View', 'Server Login', 'Server Discovery', and 'Server Polling'. Each section has specific settings that can be modified.

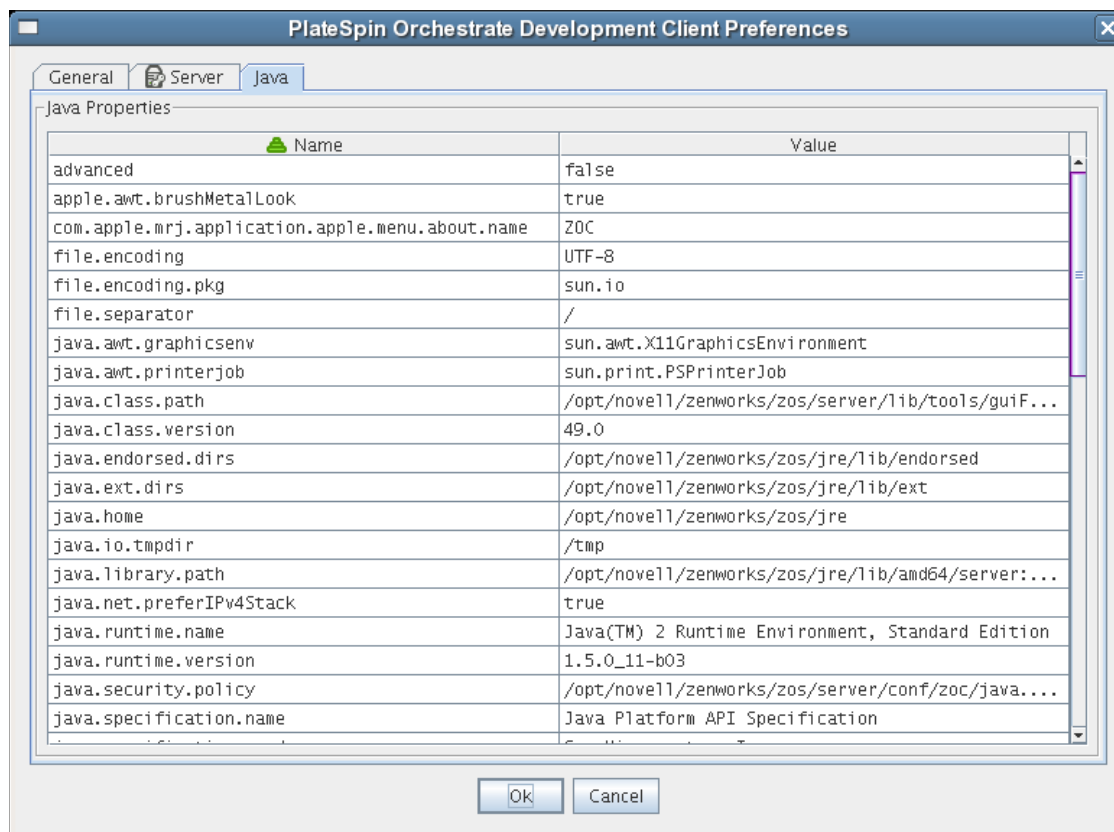
Section	Setting	Value
Log View	Maximum number of log messages	1000
	The number of buffered lines	40
	The flush log timeout (milliseconds)	1000
Server Login	Default Username	
	Default Password	*****
	Use as Default Login	<input checked="" type="checkbox"/>
	Display server log after Login	<input type="checkbox"/>
Server Discovery	Discover servers at Startup	<input checked="" type="checkbox"/>
	Discovery Timeout	3
Server Polling	Polling Interval (seconds)	1

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

Preference settings on this page that you can change are self-explanatory.

Java Properties Page

Figure 2-4 The Java Properties Page of the Orchestra Development Client Preferences



This page lists the Java property names and values that Novell uses to render the Orchestra Development Client interface in Java Swing. The list is for your information only.

2.1.3 View

The *View* menu includes various operations that let you manipulate the Orchestra Development Client display of the various PlateSpin Orchestra component views. The function of the options under this menu are self explanatory, and are a compilation of view operations that are also available from the Operations toolbar.

For more information about the View operations, see [Section 2.2, “The Orchestra Development Client Toolbar,”](#) on page 24.

2.1.4 Actions

The multiple operations listed as options under the *Actions* menu provide a quick way for you to perform operations that can also be performed (generally by right-clicking an object) in the Explorer View.

For example, if you select a *Create* option from the Actions menu, the create dialog remains open after you create each object. Here you can repeatedly create new objects in the dialog, pressing *OK* or *Create* after each is created. Similarly, in the dialog boxes of some operations in the Actions menu, you can select many objects and delete them at the same time.

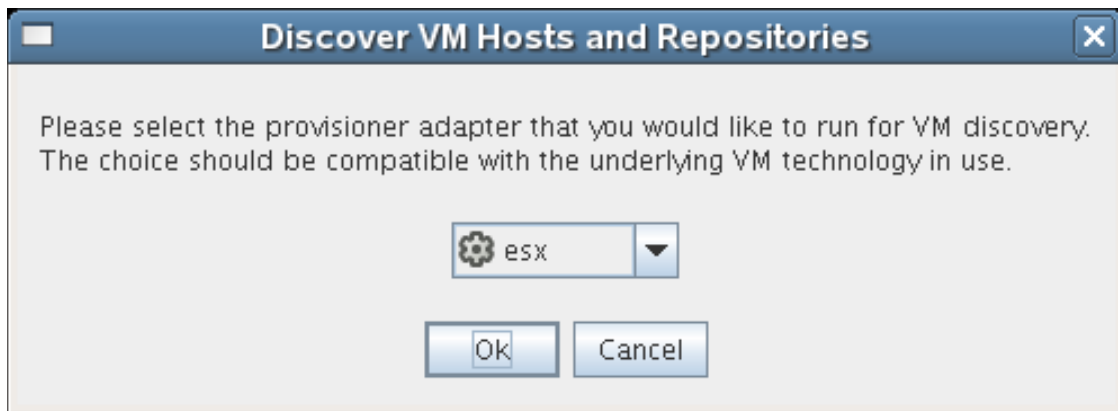
2.1.5 Provision

The *Provision* menu is added to the menu bar only if you have installed Virtual Machine Management. The multiple operations listed in the menu include two of the provisioning actions that you can execute by right-clicking a VM object in the Explorer Tree.

Discover VM Hosts & Repositories

When you select this option, the Discover VM Hosts and Repositories dialog box is displayed.

Figure 2-5 VM Discovery Dialog Box

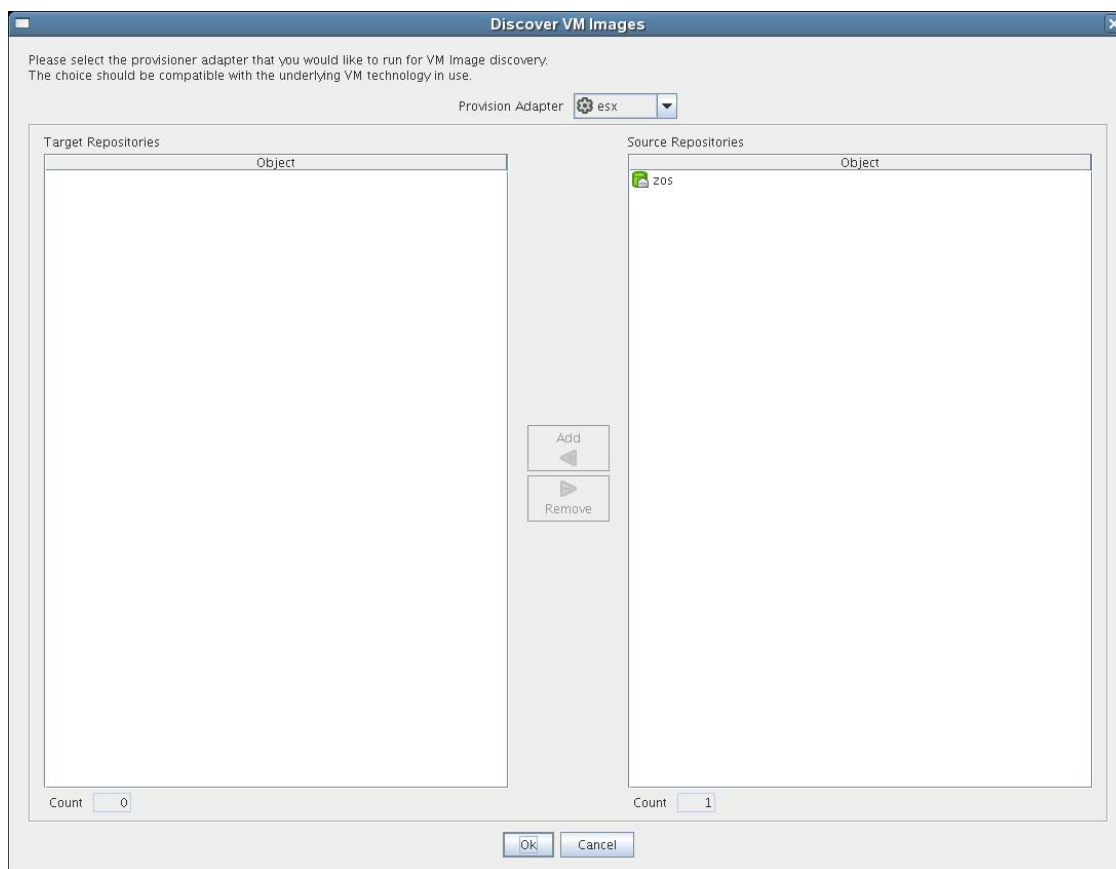


Using this dialog box, you can select a provisioning adapter (*esx*, *hyperv*, *vcenter*, *vmserver*, or *xen30*) that discovers all VM host machines where the PlateSpin Orchestrate Agent is installed and creates objects in the model. The provisioning adapter also discovers the VM Repositories where VM hosts reside.

Discover VM Images

When you select this option, the Discover VM Images dialog box is displayed.

Figure 2-6 VM Images Discover Dialog Box



Using this dialog box, you can select a provisioning adapter (esx, hyperv, vcenter, vmserver, or xen30) that discovers all VM images and creates objects in the model.

Other Provisioning Operations

The other operations listed in the menu are self-explanatory.

- ◆ Start VM Hosts
- ◆ Shutdown VM Hosts
- ◆ Shutdown VMs
- ◆ Resync VM's State
- ◆ Resync VMs's Host State
- ◆ Reset State of all VMs

2.1.6 Server

The *Server* menu lets you start a local server, log in to the server, create and display logs for logged in servers, log out from the server, and shut down a server.

- ◆ [“Select Server” on page 21](#)
- ◆ [“Discover Servers” on page 21](#)

- ♦ “Shutdown Server” on page 21
- ♦ “Login” on page 21
- ♦ “Logout” on page 21
- ♦ “Display Log” on page 22
- ♦ “Create Custom Log” on page 22

Select Server

The *Select Server* operation lets you select one of the Orchestrate Servers in your grid to log onto. When you select a server, you are required to log on. This operation accomplishes the same thing as selecting a server object from the Explorer Tree.

Discover Servers

The *Discover Servers* operation lets you launch the discovery process for servers. This is the same process that initiates (if so chosen in your server preferences) when the Orchestrate Development Client starts.

Shutdown Server

The *Shutdown Server* operation lets you shut down the current, logged on Orchestrate Server. The shutdown dialog box also lets you create a snapshot of the server state when you shut down.

Figure 2-7 The Server Shutdown Dialog Box



Login

The *Login* operation lets you establish a remote connection to another Orchestrate Server. The server IP address is required for the login. When you enter the IP address, you need to provide the username and password for the server where you are logging on.

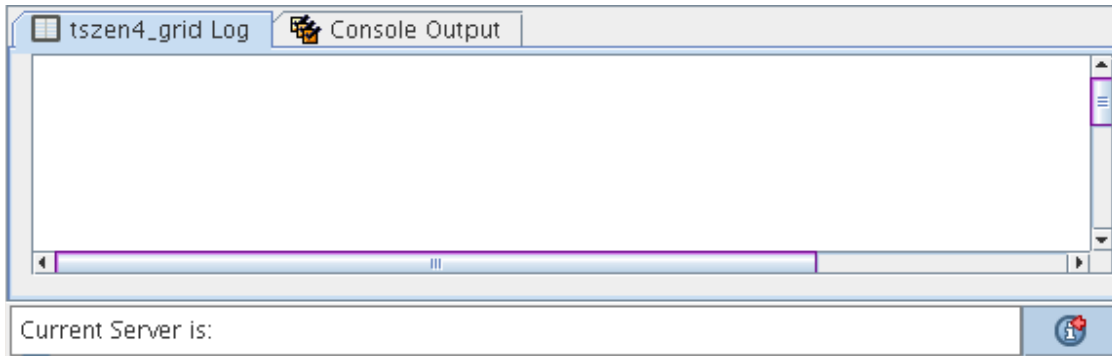
Logout

The *Logout* operation lets you log out of the current, logged on Orchestrate Server without exiting the Orchestrate Development Client. Logging out removes the server’s nodes from the Explorer Tree and its workspace views.

Display Log

The *Display Log* operation displays the default server log for the current, logged on Orchestrate Server. The display is in the Information window located at the bottom of the Orchestrate Development Client. The server log file is also located by default in the `/var/opt/novell/zenworks/zos/server/logs` directory.

Figure 2-8 Server Log Opened in Information Window of the Orchestrate Development Client



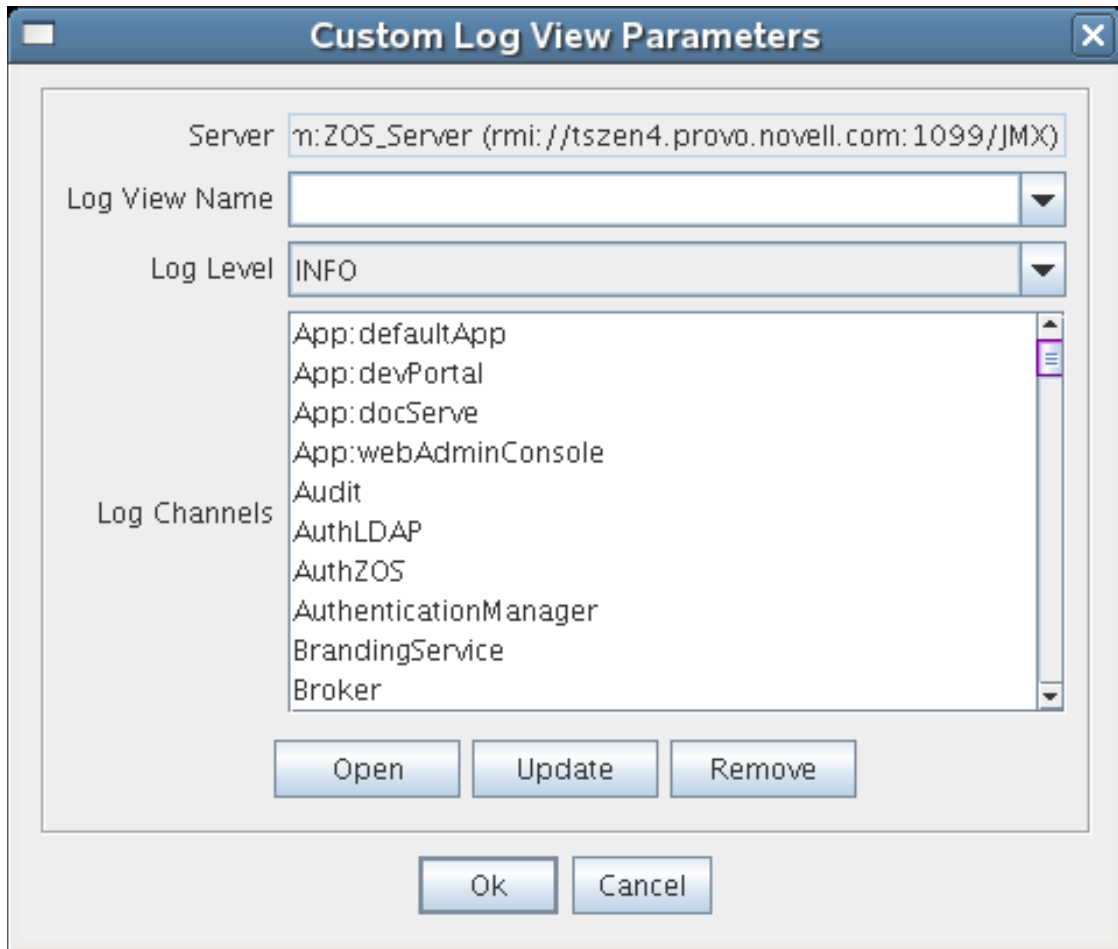
When a log is displayed, you can right-click its tab to further direct the actions of the display. You can pause logging in the window, copy the log to the clipboard, clear its contents, undock the log display as a new window, or remove it from the Information window.

If you right-click on the log display, all of the default editing capabilities of the Orchestrate Development Client are available for your use inside the window. For more information, see [Section 2.1.2, “Edit,” on page 14](#).

Create Custom Log

The *Create Custom Log* operation opens the Custom Log View Parameters dialog box.

Figure 2-9 The Custom Log View Parameters Dialog Box



By enabling a custom log, you can monitor various components of the Orchestrate Server. For example, you can view debugging information for the Audit facility. You can create, update, or remove a log view from the dialog box. You can open a custom view in the Information window by selecting *Open* in the dialog box.

Log View Name: Enter the name of the log view. This will be displayed on a tab in the log display panel.

NOTE: You can enter only alphanumeric characters and spaces in the Log View Name field.

Log Level: From the drop-down list, select the minimum log level for the log view. The log messages included in the custom view will be of this level and those of greater severity.

Log Channels: A log channel provides log information specific to an PlateSpin Orchestrate component or facility, such as the Audit facility.

When the custom view is displayed, you can right-click its tab to further direct the actions of the display. You can pause logging in the window, copy the log to the clipboard, clear its contents, undock the log display as a new window, or remove it from the Information window.

If you right-click on the log display, all of the default editing capabilities of the Orchestrate Development Client are available for your use inside the window. For more information, see [Section 2.1.2, “Edit,” on page 14](#).

2.1.7 Windows

When you right-click various views and panels in the Orchestrate Development Client, you can select the *Open in Window* option to open these views and panels in separate windows. This allows you the perspective you sometimes need when working with PlateSpin Orchestrate objects in conjunction with one another. The *Windows* menu lets you toggle between the various views or panels that are open. You can also choose to *Show All*, *Hide All*, or *Close All* of these windows.

When a given window is open, its fields and selectable dialogs remain functional so that you can perform object operations or text editing as you would when these views or panels are docked normally to the Orchestrate Development Client.




2.1.8 Help










From the *Help* menu, you can access a link to the online PlateSpin Orchestrate documentation (available in .html or .pdf format) or you can open the About box for the product, where you can view its version number, its license expiration date, and a list of its current management pack capabilities (for example, the Virtual Machine Management capability).

2.2 The Orchestrate Development Client Toolbar

The Orchestrate Development Client Toolbar includes several iconic buttons that let you perform command tasks in the Development Client workspace views and the Explorer Tree. The table below lists the functions of these buttons.

Table 2-1 *Tool Buttons from the Orchestrate Development Client Toolbar*

Tool Icon	Tool Name	Tool Function
	Back	Go back to the previous workspace view seen.
	Forward	Go forward to the next workspace view.
	Refresh	Refresh the Explorer and Workspace views.

Tool Icon	Tool Name	Tool Function
	Open/Hide Explorer	Open the Explorer Tree in a window Hide the Explorer window
		
	Cut	Cut the selected object from the workspace and copy it to the clipboard
	Copy	Copy the selected object to the clipboard while keeping the original in place
	Paste	Paste the contents of the clipboard
	Find and Replace	Open the Find dialog box
	Save	Save changes (in the workspace views or in the Explorer)
	Resource Usage Meter	(Not an active button) visual indication of resource usage. Mouse over for a listing of Active Resources, Busy resources and Available Resources, right-click to stop the meter
none (blank area)	Bookmark Toolbox	Click and drag any object from the Explorer tree into this area to create a bookmark to jump to that object's view. Right-click the bookmark to select options to open and show the object or to remove it from the toolbox. Right-click to remove all objects when some are not visible.
	Busy Indicator	(Not an active button). This pinwheel shape appears to rotate when the Server is busy performing an operation.

The Orchestrate Server and the Server Admin Objects

3

The PlateSpin Orchestrate Development Client lets you visualize the object model maintained by the Orchestrate Server and that the server uses to make resource provisioning decisions. The left pane of the Orchestrate Development Client displays a hierarchical tree known as the Explorer Tree or the Explorer View. This tree lets you navigate to different objects to see their details. Each object in the Explorer Tree is referred to as a “Grid object.” These objects can also be associated with one or more containers called Groups. When you navigate to these objects, you can edit their attributes and view more detail about their configurations.

This section includes information about the following objects that you can manage in the Explorer Tree:


- ♦ [Section 3.1, “The Orchestrate Server Object,” on page 27](#)
- ♦ [Section 3.2, “The Server Admin Object,” on page 39](#)

3.1 The Orchestrate Server Object

The highest object in the Explorer Tree is the Orchestrate Server Object, sometimes called the “grid server” object because it represents the PlateSpin Orchestrate Server acting as the holding place for all of the information used to manage objects for a single computing grid.

The PlateSpin Orchestrate Development Client is “version aware.” When the Orchestrate Development Client is launched or when server discovery is manually run, the client recognizes both current PlateSpin Orchestrate installations and old installations of discovered servers and displays their icons accordingly. This visual cue helps you to recognize when older Orchestrate Servers need to be upgraded.

Figure 3-1 Current and “Old” Server Objects



The tooltip for an Orchestrate Server lists its RMI configuration, its IP address, the directory location where the server instance was installed, and its exact version number.

The icons to the right of a current Orchestrate Server represent its policies, either those added by default upon server install and configuration, or those added later. A drop-down menu of all associated policies is opened when you right-click a policy icon. From there, you can select a policy to open in the Policy Editor. For more information about policies, see [Section 12.1, “The Policy Object,” on page 143](#).

When selected, the Server Object exposes four tabs where you can further configure its attributes. Further information about these tabs is available in the following sections:

- ♦ [Section 3.1.1, “The Orchestrate Server Info/Configuration Page,” on page 28](#)
- ♦ [Section 3.1.2, “The Orchestrate Server Authentication Page,” on page 34](#)
- ♦ [Section 3.1.3, “The Orchestrate Server Policies Page,” on page 38](#)
- ♦ [Section 3.1.4, “The Orchestrate Server Constraints/Facts Page,” on page 38](#)

3.1.1 The Orchestrate Server Info/Configuration Page

The page that opens under the *Info/Configuration* tab includes several collapsible sections on the page where you can configure the general information and attributes of the server.

- ♦ [“The Server/Cluster Panel” on page 28](#)
- ♦ [“The Data Grid Configuration Panel” on page 29](#)
- ♦ [“The Security/TLS Configuration Panel” on page 30](#)
- ♦ [“The Agent/User Session Configuration Panel” on page 31](#)
- ♦ [“The Audit Database Configuration Panel” on page 31](#)
- ♦ [“The Sentinel Server Configuration Panel” on page 32](#)
- ♦ [“The Job Limits Panel” on page 34](#)

The Server/Cluster Panel

If you are using this server in a High Availability environment, the information in this section is populated as a result of the configuration you managed during the High Availability installation. The following items are included in the section:


Server Version: A non-editable field that lists the version of this server in the form `<major>.<minor>.<point>.<build_number>`. This is the data for the fact “matrix.version”.

Is Master Server: A check box that is not selected if the server is not the Master Server in a High Availability cluster configuration.

Master Server Address: Set this value when the Orchestrate Server participates in a High Availability cluster.

External Cluster Address: Set this value when the Orchestrate Server participates in a High Availability cluster.

Cluster Addresses: Shows the hostname) or IP Addresses associated with an Orchestrate Server when it is in a High Availability configuration.

The  button opens the Attribute Element Values dialog box, where you can add, remove, or reorder addresses (element values) in an array of address choices.

For more information about using PlateSpin Orchestrate in a High Availability environment, see the [PlateSpin Orchestrate 2.5 High Availability Configuration Guide](#).

The Data Grid Configuration Panel

This section of the *Info/Configuration* tab allows for advanced configuration of datagrid related tuning parameters. The properties on the page and their descriptions are listed below.


Data Grid Root: The location of the PlateSpin Orchestrate datagrid in the file system. For example, you might change this location to use a different file system mount point (recommended when there is considerable datagrid I/O).

Cleanup Interval: The interval at which the Orchestrate Server scans User job history files on the datagrid. Job history files older than the owning user's job history retention time limit (`user.datagrid.maxhistory`) are deleted.

Cleanup Interval Enabled: Select this check box to set a flag to enable periodic job history cleanup checking. Deselect to disable the checking.

Default Multicast Rate: Sets the default data rate in bytes per second for multicast operations in which the client has not explicitly set a rate for a particular file transfer.

Max Multicast Rate: The maximum data rate(in bytes per second) that a client can specify for a multicast file transfer.

Selected Interfaces: The interfaces on which multicast file transfers are to be sent. This allows an administrator to limit multicast traffic to specific interfaces (that is, the interfaces where the agents are connected). You can add or delete interfaces by clicking the  button.

Available Interfaces: Lists the network interfaces that are available on the local machine for multicasting.

NOTE: The property is read-only and is provided for your information.

The Multicast Metrics Subpanel: This panel lets you monitor multicast data transfer, including:

- ♦ **Total Packets Sent:** The total number of multicast data packets sent by the file multicaster since the last reset of the counters.
- ♦ **Total Packets Resent:** The total number of multicast packets resent because of errors since the last counter reset.
- ♦ **Total Resend Rate:** The total packet resend rate as a percentage since the last counter reset.
- ♦ **Current Packets Sent:** The total number of multicast packets sent during the current or most recent multicast file transfer.
- ♦ **Current Packets Resent:** The total number of multicast packets resent because of errors, corruption, or loss during the current or most recent multicast file transfer.
- ♦ **Current Resend Rate:** The packet resend rate as a percentage of packets sent since the start of the current or most recent multicast file transfer.
- ♦ **Current File Size:** The file size in bytes for the current or most recent multicast file transfer.
- ♦ **Current Bytes Sent:** The number of bytes sent so far in the current or most recent multicast file transfer.
- ♦ **Current Percent Complete:** The completion percentage of the current or most recent multicast file transfer.
- ♦ **Skipped (Sparse) Bytes:** The number of bytes skipped because of long runs of zeros. These "holes" are skipped in order to reduce file transfer time for large sparse files like VM images.

- ♦ **Current Receiver Count:** The number of recipient agents for the current or most recent multicast file transfer.
- ♦ **Current File Name:** The name of the file transferred in the current or most recent multicast file transfer.

The data list includes a check box that is selected if the current multicast transfer is finished. It also includes a *Reset Stats* button that you can select to clear all of the metrics in order to begin monitoring multicast statistics from a new point in time.

The Security/TLS Configuration Panel

This section lets you configure TLS or SSL data encryption for both user and agent connections. There are four different levels of encryption that can be set for both users and nodes. These are described below. The properties in this section also let you configure the TCP/IP socket listener address and port for TLS connections.

TLS On Agent: Allows the encryption level to be set to one of four values, as described (in order of security level) below:

- ♦ **Forbid TLS for agents:** Only unencrypted connections are allowed for nodes (that is, agents) authenticating to this server. If the agent attempts to initiate encrypted communication, the connection attempt is rejected. This is the least secure of the encryption levels and is only recommended for installations where encryption is forbidden because of legal or policy restrictions, or where the performance benefits of disabling encryption outweigh security concerns.
- ♦ **Allow TLS on the agents; default to falling back to unencrypted:** Specifies that the server defaults to unencrypted communication, but the agent can optionally enable encryption.
This is the default setting for the Orchestrate Server. More secure installations might require a setting to one of the higher levels below.
- ♦ **Allow TLS on the agents; default to TLS encrypted if not configured encrypted:** The server defaults to using encryption, but the agent can optionally disable encryption.
- ♦ **Make TLS mandatory on the agents:** The Orchestrate Server rejects any connections that do not establish TLS encryption. This is the most secure encryption level because it ensures that all message communication between the node (that is, an agent) and the server is protected from tampering or interception.

TLS On Client: This setting allows the encryption level to be set to one of four values, as described (in order of security level) below.

- ♦ **Forbid TLS for clients:** Only unencrypted connections are allowed for users of this server. If the user or client attempts to initiate encrypted communication, the connection attempt is rejected. This is the least secure of the encryption levels and is only recommended for installations where encryption is forbidden because of legal or policy restrictions, or where the performance benefits of disabling encryption outweigh security concerns.
- ♦ **Allow TLS on clients; default to falling back to unencrypted:** This level specifies that the server defaults to unencrypted communication, but that the user can optionally enable encryption.

This is the default setting for the Orchestrate Server. More secure installations might require a setting to one of the higher levels below.

- ♦ **Allow TLS on agents; default to TLS encrypted if not configured encrypted:** The server defaults to using encryption, but the user can optionally disable encryption.
- ♦ **Make TLS mandatory on the clients:** The Orchestrate Server rejects any connections that do not establish TLS encryption. This is the most secure encryption level because it ensures that all message communication between the user's client programs and the server is protected from tampering or interception.

TLS Address: The port number and optional bind address for incoming encrypted connections from users and nodes. The format is `hostname:port`. For example, `10.10.10.10:8101` causes the server to accept only TLS connections on the address `10.10.10.10` on port `8101`. If `"*"` is used as the hostname, then the Orchestrate Server listens on all available network interfaces. The default is `*:8101`, which causes the Orchestrate Server to listen for encrypted sessions on all available interfaces on the system.

The Agent/User Session Configuration Panel

When nodes (agents) and users log on to the Orchestrate Server, they establish a session context that is used to manage the state of the messaging connection between client and server. This session can be revoked by the administrator, and it can also expire if the connection exceeds its maximum lifetime or idle timeout.

Agent Session Lifetime: The maximum number of seconds that an agent's session can last before the agent is disconnected and must re-authenticate with the server. A value of `-1` means "forever."

Agent Session Timeout: The idle timeout for agents. If an agent connection remains idle with no message traffic in either direction for this time period (in seconds), the session times out, and the agent is disconnected and must reauthenticate when it is ready to communicate with the server again.

Socket Keeps Agent Sessions Alive: Select this check box to set a flag that causes the server and agent to maintain a keep alive ping in order to detect hung/stalled network connections. This allows the agent to recover from hung connections and to transparently reconnect with the server.

User Session Lifetime: The maximum number of seconds that a user's session can last before the user is required to re-authenticate with the server. A value of `-1` means "forever."

User Session Timeout: The idle timeout (in seconds) for user sessions. If a user's session encounters no message traffic or requests in either direction for time, then any connection with user software is closed and the session expires. At this point, the user must re-authenticate.

Socket Keeps User Sessions Alive: Select this check box to set a flag that causes the server and user client to maintain a keep alive ping in order to detect hung/stalled network connections. This allows the agent to recover from hung connections and to transparently reconnect with the server. This setting applies only in situations where you are using custom user client software or certain subcommands of the `zos` command line utility to maintain a persistent connection.

The Audit Database Configuration Panel

This section of the Info/Configuration page lets you configure the connection to a relational database that uses a deployed JDBC driver and connection properties. The PostgreSQL driver is deployed by default.

JDBC Driver Name: Specifies the Java class for the driver.

JDBC Library: Specifies the deployed library that contains the driver.

JDBC Connection URL: Specifies the driver-dependent connection string.

Database Username: Specifies the username for database authentication.

Database Password: Specifies the password to be used for database authentication.

Is Connected: Indicates that the driver is successfully connected.

Connect (button): Click to connect through the current connection settings.

Disconnect (button): Click to disconnect the current connection.

Clear Queue (button): Clear queued records that have not yet been written to the database.

The Sentinel Server Configuration Panel

This section of the Info/Configuration page lets you configure the values needed to connect to a deployed Novell Sentinel Event Source Server, where logging events from PlateSpin Orchestrate are collected, parsed, and mapped for prioritization and subsequent administrator analysis.

For information about setting up a Sentinel Collector Server in your PlateSpin Orchestrate environment, see “[Integrating PlateSpin Orchestrate with a Sentinel Collector](#)” in the *PlateSpin Orchestrate 2.5 Installation and Configuration Guide*.

The following fields are available in the *Sentinel Server Configuration* panel:

Server Hostname: Specify the hostname of the Sentinel Event Source Server where log messages are to be sent.

Server Port Number: Specify the port number on the Sentinel Event Source Server where the PlateSpin Orchestrate Server should make its SSL connection.

Is Connected: Selected when the connection between the PlateSpin Orchestrate Server and the Sentinel Event Source Server is established.

Log Channels: Lists the log channels from which log messages are to sent to the Sentinel server.

Connect (button): Click to connect to the Sentinel Event Source Server. When the SSL connection is made, PlateSpin Orchestrate begins to send its log messages to Sentinel.

Disconnect (button): Click to disconnect the PlateSpin Orchestrate Server from the Sentinel server. When the connection ends, log messages are no longer sent to the Sentinel server.

Configure (button): Click to open the Sentinel Log Parameters dialog box. In this dialog box, you can map a log level to one or more log channels. These log channels send log messages to the Sentinel server.

For more information about PlateSpin Orchestrate log levels, see “[PlateSpin Orchestrate Log Levels Mapped to Sentinel Log Levels](#)” in the *PlateSpin Orchestrate 2.5 Installation and Configuration Guide*.

NOTE: To select multiple log channels, press Ctrl while selecting the log channel options you want. You can select only one log level at a time for mapping log channels.

The following table shows some of the log channels you can choose from and the PlateSpin Orchestrate actions that trigger sending a log message through this channel.

Table 3-1 *Log Channels and the Occasions for Sending Messages Through Each*

Log Channel Name in the Development Client (Sentinel Server Configuration Panel)	When Are Messages Sent to This Channel?
<i>ActionStatusManager</i>	<ul style="list-style-type: none"> ◆ When the status of a Grid action is updated
<i>Audit</i>	<ul style="list-style-type: none"> ◆ When the Grid interacts with the audit database
<i>AuthLDAP</i>	<ul style="list-style-type: none"> ◆ Grid-wide authentication events
<i>AuthZOS</i>	
<i>AuthenticationManager</i>	
<i>Broker</i>	<ul style="list-style-type: none"> ◆ Job execution <ul style="list-style-type: none"> ◆ start ◆ cancel
<i>Event Manager</i>	<ul style="list-style-type: none"> ◆ When a Grid object of the corresponding type is created, deleted, or its health changes to a bad state
<i>JobManager</i>	
<i>NodeManager</i>	
<i>UserManager</i>	
<i>repositoryManager</i>	
<i>vbridgeManager</i>	
<i>vdiskManager</i>	
<i>vnicManager</i>	
<i>GroupManager</i>	<ul style="list-style-type: none"> ◆ When a member is added/removed in a Group
<i>JobScheduler</i>	<ul style="list-style-type: none"> ◆ Job schedule or job trigger deployment/undeployment
<i>MBeanServer</i>	<ul style="list-style-type: none"> ◆ When internal Grid Resources are updated.
<i>PolicyManager</i>	<ul style="list-style-type: none"> ◆ Policy creation/deletion ◆ Policy association/disassociation with any Grid object
<i>Sentinel</i>	<ul style="list-style-type: none"> ◆ When the Grid interacts with a Novell Sentinel server
<i>SessionManager</i>	<ul style="list-style-type: none"> ◆ User or Resource login/logout
<i>VmManager</i>	<ul style="list-style-type: none"> ◆ Actions are performed on VMs (provision, migrate, shutdown, clone etc.). This could be initiated automatically or manually, by a user. ◆ Authorization fails during VM operation ◆ When provisioning job fails
<i>computedFact</i>	<ul style="list-style-type: none"> ◆ When computed facts are created or updated or deleted

Log Channel Name in the Development Client (Sentinel Server Configuration Panel)	When Are Messages Sent to This Channel?
<i>deployer/computedFact</i>	♦ When a corresponding resource is deployed to or undeployed from the Grid
<i>deployer/event</i>	
<i>deployer/facility</i>	
<i>deployer/jdlLibrary</i>	
<i>deployer/job</i>	
<i>deployer/library</i>	
<i>deployer/metric</i>	
<i>deployer/policy</i>	
<i>deployer/properties</i>	
<i>deployer/schedule</i>	
<i>deployer/service</i>	
<i>deployer/trigger</i>	
<i>deployer/xml</i>	

The Job Limits Panel

The facts in this section of the page are used in the default constraints to help protect the Orchestrate Server from denial-of-service attacks or badly written jobs that might otherwise get stuck in the server queue, consume resources, and cause adverse server performance.

The following fields are available in the *job.limits* panel:

max.active.jobs: Sets a global default limit on the number of active jobs.

The Orchestrate Server uses this value in the *start* constraint and does not allow more than this number of jobs (including child jobs) to be actively running at the same time. Jobs that exceed this number might be queued. See *max.queued.jobs*, below.

max.queued.jobs: Sets a global default limit on the number of queued jobs.

This value is similar to *max.active.jobs* but it is used in the *accept* constraint and limits the number of jobs sitting in a queue waiting to be started. Therefore, the maximum jobs that can be present on an Orchestrate Server is *max.active.jobs* + *max.queued.jobs*. New jobs are not be accepted by the server if they exceed this total.

job.finishing.timeout: Sets a global default limit on the timeout for job completion.

This value represents the number of seconds that the Orchestrate Server allows a job to execute its *job_cancelled_event()* (if defined) before forcibly canceling the job. This prevents jobs from potentially hanging during cancellation.

3.1.2 The Orchestrate Server Authentication Page

The *Authentication* tab opens a page with several collapsible sections where you can configure various methods for authenticating both users and resources to the PlateSpin Orchestrate Server.

- ♦ [“The Resources Panel” on page 35](#)
- ♦ [“The Users Panel” on page 35](#)
- ♦ [“The Credential Manager” on page 37](#)

The Resources Panel

The resources in a PlateSpin Orchestrate grid are actually PlateSpin Orchestrate Agents that authenticate or “register” with the PlateSpin Orchestrate Server.

Auto Register Agents: Select this check box if you want the PlateSpin Orchestrate Server to automatically register agents when they first connect to the Orchestrate Server.

The Users Panel

Only authenticated users can log into the PlateSpin Orchestrate Server. As an administrator, you can configure this authentication to use an internal user database or to externally authenticate users through an LDAP server.

Auto Register Users: Select this check box if you want the PlateSpin Orchestrate Server to automatically register users when they first connect to the Orchestrate Server.


The Enable LDAP Subpanel

Depending on the selections you make in this subpanel, the following settings are displayed:

- ♦ [“The Enable LDAP Check Box:” on page 35](#)

The Enable LDAP Check Box: Select this check box if you want the Orchestrate Server to authenticate users externally by using an LDAP server. Additional LDAP-related configuration fields are displayed when you select the check box:

Administrators

The Administrators list specifies the group names whose membership includes PlateSpin Orchestrate administrators as returned by the specified authentication provider. You can add groups to this list by clicking the  button to open an array editor dialog box, which allows groups to be added, removed, and reordered. A group must be in the format

`<provider>:<group|groupnocase>:<groupname>`, where the `<provider>` is either ZOS or LDAP. For example, adding `LDAP:groupnocase:XyZ` allows users reported by the LDAP server as members of a group xyz, or XYZ, xYz, etc. to authenticate as an administrator. To enforce to case-sensitive matching, use `LDAP:group:XyZ` instead. Non-case-sensitive matching is needed for Active Directory servers.

Active Directory Service Settings

If you select Active Directory Service in the *Server Type* drop down list, the following settings are available:


Directory Name: The name of the Active Directory Service server.

Servers: A list of strings containing `server:port` entries for a list of servers to be used.

Each entry can be of one of three forms:

- ♦ `<hostname>`
- ♦ `<hostname>:<port>`
- ♦ `<hostname>:<port>:<sslport>`

In all cases, <hostname> is a resolvable DNS name or an IP address. If SSL or TLS is in use, the hostname must exactly match the name on the ADS server SSL certificate.

You can modify this list by clicking the  button to open an Attribute Element Values dialog box, where you can add, remove, or change the order of server names.

Advanced: The settings in this section are for more selective ADS authentication.

- ♦ **SSL:** If the accompanying *Start TLS* check box is not selected and if the ADS server's SSL certificate has been installed on the PlateSpin Orchestrate Server JVM, this option securely connects to the ADS server through SSL encryption.

The older LDAP protocol (`ldaps://`) is used for the connection.

- ♦ **Start TLS:** Selecting this option immediately promotes the connection to SSL encryption by bypassing the older protocol in favor of the LDAPv3 Start TLS extended operation on the non-SSL LDAP port. To use this option, the ADS server's SSL certificate must be installed on the JVM of the PlateSpin Orchestrate Server.
- ♦ **Query Account:** The account name that is to be used for querying group information on authenticated users.
- ♦ **Query Password:** The clear text password used to authenticate the query account on the LDAP server.

Generic Settings

When you select *Generic LDAP Directory Service* as the Server Type, the following additional settings are displayed:

Base Domain Name: The Root DN of the LDAP server's directory tree. This must be obtained by the administrator, and is usually in the form of `dc=adsroot,dc=novell,dc=com`.

User Attribute: The attribute on a user's entry that identifies his or her login account name. For ADS servers, this attribute is `sAMAccountName`.

User Filter: The name of the filter to be used in the lookup for the user's LDAP distinguished name.

For ADS, this prefix is `cn=Users`.

User Prefix: The prefix used to define the LDAP subtree within the BaseDN tree that contains user accounts. If you leave this property blank, the Orchestrate Server uses the BaseDN.

Group Attribute: Specifies the attribute of a group entry describing the login name of that group.

Group Filter: A filter to be used in the lookup for group memberships on some LDAP schemas. The filter can use either `${USER_NAME}` or `${USER_DN}` to substitute that value. For example: `memberUid=${USER_NAME}`.

Group Prefix: The prefix used to define the LDAP subtree within the BaseDN tree that contains group accounts.

Not used for Active Directory authentication.

Group DNA Attribute: The directory root where all queries for a user's group memberships (stored as a list of "member of" attributes on the user's entry on an ADS server) are to occur.

Nested DNA Attribute: The attribute of a group entry where subgroups can be queried.

The Credential Manager

As a data center administrator, you often have to provide credentials and certificates as you interact with the different hypervisor technologies– the Amazon EC2 or vSphere technologies, in particular. PlateSpin Orchestrate lets you store this data in a centralized, secure (no cleartext passwords are accessible) location in its Credential Manager.

NOTE: PlateSpin Orchestrate uses TripleDES password-based encryption in its Credential Manager to encrypt stored credentials and certificates.

The Credential Manager, located in the Authentication page of the Orchestrate Server Grid object in the Development Client, includes the following sections:

- ♦ [“The Stored Credentials Panel” on page 37](#)
- ♦ [“The Stored Certificates Panel” on page 37](#)

The Stored Credentials Panel

The Stored Credentials panel displays a list of names of credential sets that you have created. You can create additional credentials if you select *Add Credential* and fill in the following fields:

Name: (Required) The name that you want to use to refer to this credential set.

User: (Required) The username with rights to administer objects in this grid.

Secret/Password: (Required) The password that authenticates the user.

Type: (Optional) A user-defined string that lets similar credentials be put into a category or group. For example, you might have a “type” of credential for the amazon-ec2 provisioning adapter and another type for the vsphere provisioning adapter.

Stored Credentials Password: (Conditional) If you want to change the password element of your stored credentials, click *Change* and enter the new password.

This password is stored as a fact on the Matrix grid object. In the Fact Editor, the fact is listed as `matrix.credential.manager.passphrase`. It is used to encrypt the stored passwords. By default the password is `CHANGE_THIS_PASSWORD`. We recommend that you select a new password to use for encrypting stored passwords.

The Stored Certificates Panel

In order to trust certificates not signed by well known certificate authorities, PlateSpin Orchestrate lets you store certificates that are trusted by Java.

NOTE: Public/Private key pairs can be stored as certificates. This is useful if you need to manage amazon- ec2 key pairs.

The Stored Certificates panel displays a list of stored certificates. These certificates are not mapped to anything other than the name or identifier that you assign. They are not stored in a trust store, but their PEM-encoded representation is encrypted and stored alongside the credentials referred to above. Trust stores are generated on demand and are available to the Orchestrate Agents.

You can create additional trust stores if you select *Add Certificate* and fill in the following fields:

Currently (Orchestrate version 2.5), this functionality is used only by the Orchestrate vsphere provisioning adapter.

Identifier: (Required) The name that you want to use to refer to this trust store.

Location: (Required) Where the certificate should be obtained. This can be either a file (which you can browse to find on the local machine), or as an HTTPS server.

Select *Browse* if you want to select an existing a PEM-encoded certificate file from the local machine.

If you want to provide the actual URL for the certificate, open the drop-down list, select *HTTPS*, then enter the URL. The HTTPS server address can be entered as:

```
https://your.server.name
```

or as

```
your.server.name
```

or as

```
https://your.server.name:<sslport>
```

With this address, Orchestrate retrieves the public server certificate from the server and then stores it in a secure location.

Group: (Optional) A user-defined string used for grouping related certificates. For example, you might have a grouping called “vsphere” when you are managing resources in a multiple-vSphere Server environment.

3.1.3 The Orchestrate Server Policies Page

The Policies tab opens a page that contains a policy viewer for each of the policies associated with the Server Object.

NOTE: You can edit a policy by right-clicking a policy icon, selecting *Edit Policy*, and clicking the save icon.

3.1.4 The Orchestrate Server Constraints/Facts Page

The *Constraints/Facts* tab opens a page that shows all of the effective constraints and facts for the Server object. The Server object has an associated set of facts and constraints that define its properties. By building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any system resource by managing an object’s facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode *r/o* have read-only values, which can be viewed by using the pencil icon, but changes cannot be made.

3.2 The Server Admin Object

The Server Admin object lists the accessible PlateSpin Orchestrate Servers and their deployed components. Clicking a deployed component displays information about that component's associated Deployment Session.

The Job Object

A job is deployed to the Orchestrate Server to automate processes, such as coordinating VM provisioning, high-performance computing, or general data center management. Jobs consist of Job Development Language (JDL) scripts and might have one or more policies associated with them. Policies define job arguments and other facts that are used by the job.

Usually a job has logic that runs on the PlateSpin Orchestrate Server itself and schedules work to run on one or more managed resources that are running the PlateSpin Orchestrate Agent. The logic that is dispatched and run on the managed resources is called a joblet. A job might or might not define one or more joblets.

A JDL script is partitioned into a Job section and one or more Joblet sections. The joblet sections of the script describe most of the work of a job. The PlateSpin Orchestrate Server dispatches joblets to resources in the grid where the work is done.

The Job object also contains facts with attributes that are used for job and joblet control. Policies associated with the job also control the job. The Orchestrate Development Client has an administrative (“admin”) view in the Explorer Panel that lets you edit these objects.

This section includes information about a Job object that is visible in the Explorer view and the accompanying Admin view of the Orchestrate Development Client:

- ♦ [Section 4.1, “Job Groups,” on page 41](#)
- ♦ [Section 4.2, “The Job Info/Groups Tab,” on page 41](#)
- ♦ [Section 4.3, “The JDL Editor Tab,” on page 51](#)
- ♦ [Section 4.4, “The Job Library Editor Tab,” on page 52](#)
- ♦ [Section 4.5, “The Job Policies Tab,” on page 53](#)
- ♦ [Section 4.6, “The Job Constraints/Facts Tab,” on page 54](#)



4.1 Job Groups

Any group object displayed in the Explorer panel represents a collection of similar object types. Groups can also be created automatically, such as when a provisioning adapter (PA) discovers a local repository on a VM host. For example, the xen30 PA, upon discovery of a VM host, automatically creates a local repository for that VM host and places the created repository in a xen30 repository group. You can also create groups manually in the Development Client, either by clicking the *Actions* menu and choosing *Create Job Group* or by right-clicking a Job Group object (anywhere in the Job hierarchy) and selecting *New Job Group*.

4.2 The Job Info/Groups Tab

The page that opens under the *Info/Configuration* tab of the Job admin view includes several collapsible sections on the page where you can configure the general information and attributes of the job.

- ♦ [Section 4.2.1, “Info,” on page 42](#)
- ♦ [Section 4.2.2, “Groups,” on page 51](#)

NOTE: Whenever you make changes to any Grid object, the write icon is superimposed on the objects icon , signifying that the object has been altered. If you want to save the changes you have made, you need to click the save icon  on the Development Client toolbar.

4.2.1 Info

The following fields on the Information panel provide facts for the Job object:

- ♦ [“Show Inherited Fact Values Check Box” on page 42](#)
- ♦ [“Job Control Settings” on page 42](#)
- ♦ [“Joblet Control Settings” on page 46](#)
- ♦ [“Automatic Resource Provisioning Settings” on page 48](#)
- ♦ [“Resource Preemption Settings” on page 48](#)
- ♦ [“Job Counts” on page 49](#)
- ♦ [“Job History” on page 49](#)

Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached or inherited policies. These fact values are read only (non-editable).

Job Control Settings

The Job Control Settings panel on the Info/Groups page includes the following fields:

NOTE: Tooltip text is available when you mouse over any of these fields.

Description: Enter information in this box that describes the nature or purpose of this job.

In the Fact Editor, this fact is listed as `job.description`:

```
<fact name="job.description" value="" type="String" />
```

Enabled: This check box is selected by default. When it is selected (it has a value of true), the job is enabled (and it is ready to run).

In the Fact Editor, this fact is listed as `job.enabled`:

```
<fact name="job.enabled" value="true" type="Boolean" />
```

Job Visible to Users: This check box is selected by default. When it is selected (it has a value of true), the job can be viewed in the Development Client, through the use of command line queries, or in the Orchestrate Server Portal. Deselecting this check box does not keep the job from running.

In the Fact Editor, this fact is listed as `job.visible`:

```
<fact name="job.visible" value="true" type="Boolean" />
```

JDL Debug Tracing: This check box is not selected by default. When it is selected (it has a value of true), the job log includes tracing information when job events are executed.

In the Fact Editor, this fact is listed as `job.tracing`:

```
<fact name="job.tracing" value="false" type="Boolean" />
```

Job Type: Lets you choose the job type that applies to this job. This setting is optional and is leveraged by the server in order to provide better quality completion time calculation for the job.

The job type options (completion time algorithms) include:

- ♦ **normal:** The default job type. If this job has joblets, the job is based on PSPACE progression algorithm. If it does not have joblets, it is based on historical wall time average.
- ♦ **workflow:** This job type does not offer a time algorithm to the server.
- ♦ **pspace:** If this job has joblets, the job is based on PSPACE progression. If it does not have joblets, do not offer a time algorithm.
- ♦ **fixedtime:** This job type directs the server to use a time algorithm based on historical wall time average.
- ♦ **fixedgcycles:** If this job has joblets, the job is based on average gcycles and current consumption rate. If it does not have joblets, the job is based on historical wall time average.

NOTE: You can change this setting at runtime to refine the calculation time as the job progresses. For example, the *zosmake* job might start out as type *normal*, but when all tasks have been submitted, you could change it to type *workflow* to allow its subjobs to drive the end time.

In the Fact Editor, the Job Type fact is listed as `job.jobtype`:

```
<fact name="job.jobtype" value="normal" type="String" />
```

Job Timeout: The amount of time (in seconds) after which the server can take action to cancel the whole job, including all joblets and subjobs. A value of `-1` indicates no timeout.

In the Fact Editor, this fact is listed as `job.timeout`:

```
<fact name="job.timeout" value="-1" type="Integer" />
```

Job Auto Terminate: This check box is selected by default. When it is selected (it has a value of `true`), the job ends when all child jobs and joblets are executed.

In the Fact Editor, this fact is listed as `job.autoterminate`.


```
<fact name="job.autoterminate" value="true" type="Boolean" />
```

Provision Adapter Hook Jobs: The name of a job that implements administrator-defined pre- or post-provisioning hooks.

NOTE: This fact is visible in the *Info/Groups* tab only when a provision adapter job is selected.

In the Fact Editor, this fact is listed as an array:

```
<fact name="job.paHooksVmJob">
  <array type="String">
    </array>
</fact>
```

You can edit this array by clicking the  button to open the *Attribute Element Values* dialog box, where you can add or remove fact specifications to the array of element choices. For more information, see “[Provisioning Adapter Hooks](#)” in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*.

Queue Type: Lets you choose the queue type that applies to this job. This setting is optional and is leveraged by the server to provide a better start time calculation for the job.

The queue type options (start time algorithms) include:

- ♦ **none:** The start time is always unknown for jobs that are queued.
- ♦ **pfifo:** Packet First In First Out. The start time implemented through policies. The server is directed to look at the job as having a finite number of active slots, so its start time depends on its position in the queue and the estimated end time of running jobs of this type. The FIFO queue for this queue reshuffles based on priority.
- ♦ **fifo:** First In First Out. The start time implemented through policies. The server is directed to look at the job as having a finite number of active slots, so its start time depends on its position in the queue (first-come, first-served) and the estimated end time of running jobs of this type. The FIFO queue for this job does not reshuffle based on priority.
- ♦ **lifo:** Last In First Out. The start time is implemented through policies. The server is directed to look at the job as having a finite number of active slots, so its start time depends on its position in the queue and the estimated end time of running jobs of this type. The queue for this job does not reshuffle based on priority.
- ♦ **fixedtime:** The start time is based on the historical average queue time. This can be explicitly overridden through setting the `job.history.queueTime.average` fact.

In the Fact Editor, this fact is listed as `job.queueType`:

```
<fact name="job.queueType" value="pfifo" type="String" />
```

Job Queued Timeout: The amount of time (in seconds) after which the server can take action to cancel a queued job, including all joblets and subjobs. A value of -1 indicates no timeout.

In the Fact Editor, this fact is listed as `job.queuedTimeout`:

```
<fact name="job.queuedTimeout" value="-1" type="Integer" />
```

Resource Match Cache TTL: Specifies the job's willingness to allow resource matches to be cached if the Job Scheduler becomes too loaded. The value is the time (in seconds) to live (TTL) of the cache. Enter a value less than zero (<0) to disable caching.

In the Fact Editor, this fact is listed as `job.cacheResourceMatches.ttl`:

```
<fact name="job.cacheResourceMatches.ttl" value="30" type="Integer" />
```

Preemptible: This check box is not selected by default. When it is selected (it has a value of true), you set the job's ability to be preempted. This setting can be overridden by the job instance.

In the Fact Editor, this fact is listed as `job.preemptible`:

```
<fact name="job.preemptible" value="false" type="Boolean" />
```

Restartable: This check box is not selected by default. When it is selected (it has a value of true), you set the job's ability to be restarted when the server restarts. This setting can be overridden by the job instance.

In the Fact Editor, this fact is listed as `job.restartable`:

```
<fact name="job.restartable" value="false" type="Boolean" />
```

Absolute Max Joblets: Specifies the absolute maximum number of joblets that you want this job to schedule.

In the Fact Editor, this fact is listed as `job.joblet.max`:

```
<fact name="job.joblet.max" value="1000" type="Integer" />
```

Max Joblet Failures: Specifies the number of non-fatal joblet errors that you want this job to tolerate before the job fails completely. Set the value at -1 to attempt to continue after errors.

In the Fact Editor, this fact is listed as `job.joblet.maxfailures`:

```
<fact name="job.joblet.maxfailures" value="0" type="Integer" />
```

Max Node Failures: Specifies the number resource failures that you want this job to tolerate before the node is excluded from further joblet processing. Set the value at -1 to specify that limited failures are acceptable.

In the Fact Editor, this fact is listed as `job.maxnodefailures`:

```
<fact name="job.maxnodefailures" value="2" type="Integer" />
```

Max Resources: Specifies the absolute maximum number of resources that you want the job to use at one time. PlateSpin Orchestrate does not exceed the value set here. Set the value at -1 to specify unlimited resources.

In the Fact Editor, this fact is listed as `job.maxresources`:

```
<fact name="job.maxresources" value="-1" type="Integer" />
```

Max Joblets Running: Specifies the absolute maximum number of joblets that you want the job to have running at one time. PlateSpin Orchestrate does not exceed the value set here. Set the value at -1 to specify unlimited joblets.

In the Fact Editor, this fact is listed as `job.joblet.maxrunning`:

```
<fact name="job.joblet.maxrunning" value="-1" type="Integer" />
```

Max Joblets Per Resource: Specifies the absolute maximum number of joblets that you want the job to occupy on a resource. Set the value at -1 to specify unlimited joblets.


In the Fact Editor, this fact is listed as `job.joblet.maxperresource`:

```
<fact name="job.joblet.maxperresource" value="-1" type="Integer" />
```

Resource Selection Ranking: Displays ranking specification used to select suitable resources. The element syntax is `fact/order` where `order` is either ascending or descending

In the Fact Editor, this fact is listed as an array:

```
<fact name="job.resources.rankby">
  <array>
    <string>resource.loadaverage/a</string>
    <string>resource.anything/a</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open the *Attribute Element Values* dialog box, where you can add or remove fact specifications to the array of element choices.

Persist Facts on Completion: This check box is not selected by default. When it is selected (it has a value of true), you specify that the Grid objects that this job modifies are persistent after the job. This setting is available and applicable only in a High Availability setup.

In the Fact Editor, this fact is listed as `job.persistfactsonfinish`:

```
<fact name="job.persistfactsonfinish" value="false" type="Boolean" />
```

Joblet Control Settings

Joblet Timeout: Specifies the amount of time (in seconds) you want the Orchestrate Server to wait until cancelling the joblet. Set the value at -1 to specify no timeout.

In the Fact Editor, this fact is listed as `job.joblet.timeout`:

```
<fact name="job.joblet.timeout" value="-1" type="Integer" />
```

Max Joblet Retries: Specifies the number of joblet retries (of any type) to be attempted before the Orchestrate Server considers the joblet as failed. A value of zero (0) specifies that the joblet should not be retried. A value of less than zero (<0) specifies that the joblet should be continually retried.

In the Fact Editor, this fact is listed as `job.joblet.maxretry`:

```
<fact name="job.joblet.maxretry" value="0" type="Integer" />
```

Retry Limit (Forced): Specifies the number of forced joblet retries (requested by the joblet to run on another resource) to be allowed before the Orchestrate Server considers the joblet as failed. A value of zero (0) specifies that the joblet should not be retried. A value of less than zero (<0) specifies that the joblet should be continually retried. This value should never exceed the value in `job.joblet.maxretry`.

In the Fact Editor, this fact is listed as `job.joblet.retrylimit.forced`:

```
<fact name="job.joblet.retrylimit.forced" value="-1" type="Integer" />
```

Retry Limit (Unforced): Specifies the number of unforced joblet retries to be allowed before the Orchestrate Server considers the joblet as failed. A value of zero (0) specifies that the joblet should not be retried. A value of less than zero (<0) specifies the joblet should be continually retried. This value should never exceed the value in `job.joblet.maxretry`.

In the Fact Editor, this fact is listed as `job.joblet.retrylimit.unforced`:

```
<fact name="job.joblet.retrylimit.unforced" value="-1" type="Integer" />
```

Retry Limit (Resource Disconnect): Specifies the number of joblet retries caused by unexpected resource disconnect to be allowed before the Orchestrate Server considers the joblet as failed. A value of zero (0) specifies that the joblet should not be retried. A value of less than zero (<0) specifies the joblet should be continually retried. This value should never exceed the value in `job.joblet.maxretry`.

In the Fact Editor, this fact is listed as `job.joblet.retrylimit.disconnect`:

```
<fact name="job.joblet.retrylimit.disconnect" value="-1" type="Integer" />
```

Retry Limit (Timeout): Specifies the number of joblet retries caused by server-initiated joblet timeout to be allowed before the Orchestrate Server considers the joblet as failed. A value of zero (0) specifies that the joblet should not be retried. A value of less than zero (<0) specifies that the joblet should be continually retried. This value should never exceed the value in `job.joblet.maxretry`.

In the Fact Editor, this fact is listed as `job.joblet.retrylimit.timeout`:

```
<fact name="job.joblet.retrylimit.timeout" value="-1" type="Integer" />
```

Immediately Retry Failed Joblet: This check box is not selected by default. When it is selected (it has a value of true), you specify that you want the system to immediately retry a joblet rather than waiting until all others are either running or complete before retrying.

In the Fact Editor, this fact is listed as `job.joblet.immediateretry`:

```
<fact name="job.joblet.immediateretry" value="true" type="Boolean" />
```

Max Joblet Wait Time: Specifies the amount of time (in seconds) you want a resource to wait before being utilized by a joblet. A setting of -1 indicates no timeout.

In the Fact Editor, this fact is listed as `job.joblet.maxwaittime`:

```
<fact name="job.joblet.maxwaittime" value="-1" type="Integer" />
```

Joblet JDL Debug Tracing: This check box is not selected by default. When it is selected (it has a value of true), you specify that you want the joblet to include tracing information on the job log as it executes joblet events.

In the Fact Editor, this fact is listed as `job.joblet.tracing`:

```
<fact name="job.joblet.tracing" value="false" type="Boolean" />
```

Joblet Run Type: Use this list to select whether or not the file and executable operations that run in the joblet are in behalf of the job user.

- ♦ **RunAsJobUserFallingBackToNodeUser:** (The default setting.) If this option is selected, any joblet logic executes as the local user with the same name as the grid user. If a local user of a matching name is not available, the joblet logic runs as the same user who is running the Orchestrate Agent (also known as the Node User). By default, the agent (Node User) is `root`.
- ♦ **RunOnlyAsJobUser:** If this option is selected, any joblet logic executes as the local user using the same name as the grid user (that is the Orchestrate Server user who matches the PlateSpin Orchestrate username). If a local user of a matching name is not available, the joblet logic (and perhaps the job) fails. By default, the agent (Node User) is `root`.
- ♦ **RunOnlyAsNodeUser:** If this option is selected, any joblet logic runs as the same user who is running the Orchestrate Agent (also known as the “Node User”). It does not run as the OS user whose username matches the PlateSpin Orchestrate user name. By default, the agent (Node User) is `root`.

In the Fact Editor, this fact is listed as `job.joblet.runtype`:

```
<fact name="job.joblet.runtype" value="RunAsJobUserFallingBackToNodeUser" type="String" />
```

Automatic Resource Provisioning Settings

Max Resource Provisions: specifies the number of resources that can be automatically provisioned in behalf of this job. A setting of zero (0) turns off automatic provisioning behavior. A setting of -1 allows unlimited provisioning.

In the Fact Editor, this fact is listed as `job.provision.maxcount`:

```
<fact name="job.provision.maxcount" value="0" type="Integer" />
```

Max Pending Provisions: Specifies the number of resources that can be automatically provisioned at one time (that is, simultaneously) in behalf of this job. A setting of less than or equal to zero (≤ 0) turns off automatic provisioning behavior.

In the Fact Editor, this fact is listed as `job.provision.maxpending`:

```
<fact name="job.provision.maxpending" value="1" type="Integer" />
```

Max Resource Provision Failures: Specifies the maximum number of provision failures resources to be tolerated before excluding the node from future automatic provisioning. A setting of -1 indicates that unlimited failures are acceptable.


In the Fact Editor, this fact is listed as `job.provision.maxnodefailures`:

```
<fact name="job.provision.maxnodefailures" value="1" type="Integer" />
```

Provision Selection Ranking: Displays ranking the specification used to select suitable resources to automatically provision. The element syntax is `fact/order` where order is either ascending or descending.

In the Fact Editor, this fact is listed as an array:

```
<fact name="job.provision.rankby">
  <array type="String">
    </array>
</fact>
```

You can edit this array by clicking the  button to open the *Attribute Element Values* dialog box, where you can add or remove fact specifications for the array of element choices.

Host Selection Strategy: Lets you choose the type of strategy you want to use in finding a host for any automatically provisioned resource. The choices include:

- ♦ **queue:** Directs the server to use the default affinity wait period defined by the resource before considering all possible hosts. The request is queued until a suitable resource becomes available or a requesting job completes.
- ♦ **immediate:** Directs the server to immediately consider the affinity host before trying to find any matching resources and to fail if a suitable resource is not available.

In the Fact Editor, this fact is listed as `job.provision.hostselection`:


```
<fact name="job.provision.hostselection" value="immediate" type="String" />
```

Resource Preemption Settings

Job Selection Ranking: Displays the ranking specification used to select suitable jobs to automatically preempt on a resource. Element syntax is `fact/order` where order is either ascending or descending.

In the Fact Editor, this fact is listed as an array:

```
<fact name="job.preemption.rankby">
  <array>
    <string>jobinstance.priority/a</string>
    <string>jobinstance.joblets.running/d</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open the *Attribute Element Values* dialog box, where you can add or remove fact specifications for the array of element choices.

Job Counts

Total Instances: Displays the total number of job instances of this type that exist in the PlateSpin Orchestrate system.

In the Fact Editor, this fact is listed as `job.instances.total`:

```
<fact name="job.instances.total" value="0" type="Integer" />
```

Active Instances: Displays the total number of job instances of this type that are in a queued state in the PlateSpin Orchestrate system.

In the Fact Editor, this fact is listed as `job.instances.active`:

```
<fact name="job.instances.active" value="0" type="Integer" />
```

Queued Instances: Displays the total number of job instances of this type that are active in the PlateSpin Orchestrate system.

In the Fact Editor, this fact is listed as `job.instances.queued`:

```
<fact name="job.instances.queued" value="0" type="Integer" />
```

Job Accounting Group: Lets you select the Job Group whose statistics are updated by default when the job runs.

In the Fact Editor, this fact is listed as `job.accountinggroup`:

```
<fact name="job.accountinggroup" value="all" type="String" />
```

Job Resource Group: Lets you select the default Resource Group whose members and any of its resource policies are selected for this job.

In the Fact Editor, this fact is listed as `job.resourcegroup`:

```
<fact name="job.resourcegroup" value="all" type="String" />
```

Job History

Shared Instance Count: (Read only) Displays the total number of job instances (including those denied by “accept” constraints) of this job that have ever been initiated on this PlateSpin Orchestrate system.

In the Fact Editor, this fact is listed as `job.history.jobcount`:

```
<fact name="job.history.jobcount" value="0" type="Integer" />
```

Completed Count: (Read only) Displays the total number of job instances (including those denied by “accept” constraints) of this job that have been canceled.

In the Fact Editor, this fact is listed as `job.history.jobcount.complete`:

```
<fact name="job.history.jobcount.complete" value="0" type="Integer" />
```

Cancelled Count: (Read only) Displays the total number of job instances (including those denied by “accept” constraints) of this job that have been completed.

In the Fact Editor, this fact is listed as `job.history.jobcount.cancelled`:

```
<fact name="job.history.jobcount.cancelled" value="0" type="Integer" />
```

Failed Count: (Read only) displays the total number of job instances of this type that have failed.

In the Fact Editor, this fact is listed as `job.history.jobcount.failed`:

```
<fact name="job.history.jobcount.failed" value="0" type="Integer" />
```

Total Cost: Displays the total cost of running this job. The amount is calculated since the job was deployed or last modified.

In the Fact Editor, this fact is listed as `job.history.cost.total`:

```
<fact name="job.history.cost.total" value="0.0000" type="Real" />
```

Average Cost: displays the average cost of running this job. The amount is calculated since the job was deployed or last modified and is updated only if the job finishes successfully.

In the Fact Editor, this fact is listed as `job.history.cost.average`:

```
<fact name="job.history.gcycles.average" value="0" type="Integer" />
```

Total Runtime: Displays the total runtime (in seconds) since the job was deployed.

In the Fact Editor, this fact is listed as `job.history.runtime.total`:

```
<fact name="job.history.runtime.total" value="0" type="Integer" />
```

Average Runtime: Displays the average runtime (in seconds) since the job was deployed.

In the Fact Editor, this fact is listed as `job.history.runtime.average`:

```
<fact name="job.history.runtime.average" value="0" type="Integer" />
```

Total Execution Time: Displays the total combined resource wall time (in seconds) of all work performed on behalf of this job since the job was deployed.

In the Fact Editor, this fact is listed as `job.history.time.total`:

```
<fact name="job.history.time.total" value="0" type="Integer" />
```

Average Execution Time: Displays the average resource wall time (in seconds) of all work performed on behalf of this job since the job was deployed.

In the Fact Editor, this fact is listed as `job.history.time.average`:

```
<fact name="job.history.time.average" value="0" type="Integer" />
```

Total Grid Time: Displays the total amount of normalized grid time (in gcycles) consumed by this job since deployment.

In the Fact Editor, this fact is listed as `job.history.gcycles.total`:

```
<fact name="job.history.gcycles.total" value="0" type="Integer" />
```

NOTE: A gcycle can be thought of as a normalized second of compute time. It is a relative measure that is approximately a second of the real processing time of a 2Ghz Pentium* class Intel* processor.

Average Grid Time: Displays the average amount of normalized grid time (in gcycles, which is a normalized grid cycle) consumed by running this job. The value is updated only if the job finishes successfully.

In the Fact Editor, this fact is listed as `job.history.gcycles.average`:

```
<fact name="job.history.gcycles.average" value="0" type="Integer" />
```

Total Queue Time: Displays the total amount of time (in seconds) since deployment that the job has spent in a queued state.

In the Fact Editor, this fact is listed as `job.history.queueetime.total`:

```
<fact name="job.history.queueetime.total" value="0" type="Integer" />
```

Average Queue Time: Displays the average amount of wall time (in seconds) spent waiting for this job to start.

In the Fact Editor, this fact is listed as `job.history.queueetime.average`:

```
<fact name="job.history.queueetime.average" value="0" type="Integer" />
```

Average Sample Size: Displays the total number of points you want to use in the trailing average calculation for all historical averages.

In the Fact Editor, this fact is listed as `job.history.samplesize`:

```
<fact name="job.history.samplesize" value="2" type="Integer" />
```

NOTE: A trailing average is the mean average measured over the last *x* datapoints.

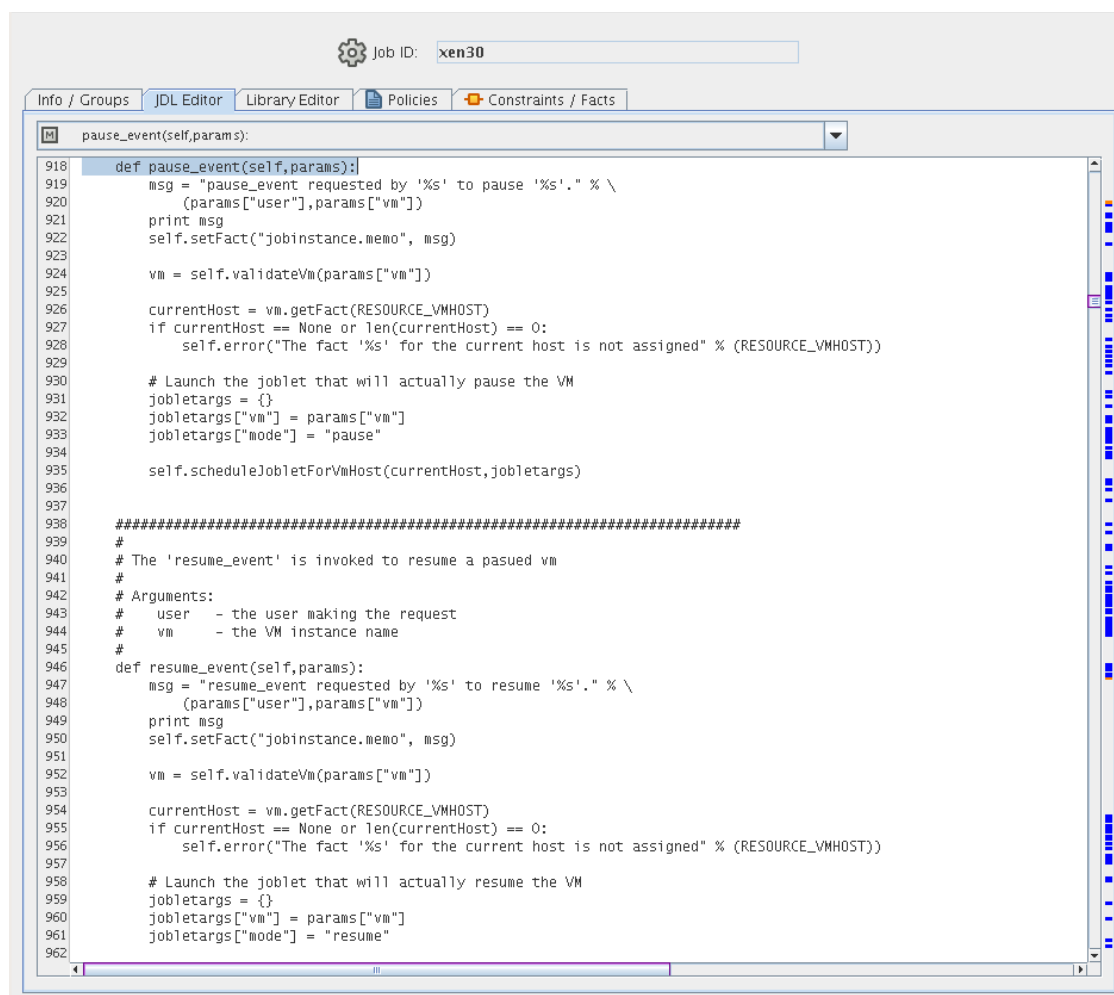
4.2.2 Groups

This section of the Info/Groups page lists the groups of Job objects in the grid. Click *Choose* to open the Job Group Selection dialog box. In this dialog box, you can choose which Job Groups to display in the Explorer Panel by selecting a group and then clicking *Add* or *Remove* to move it to or from the *Source Job Groups* list.

4.3 The JDL Editor Tab

The *JDL Editor* tab of the Job admin view opens an editor where you can inspect and modify the Job Description Language (JDL) code. This code consists of a Python script that contains the bits to control a job. The JDL code for each job includes commented documentation to explain the job's purpose and methods for implementation.

Figure 4-1 The JDL Editor

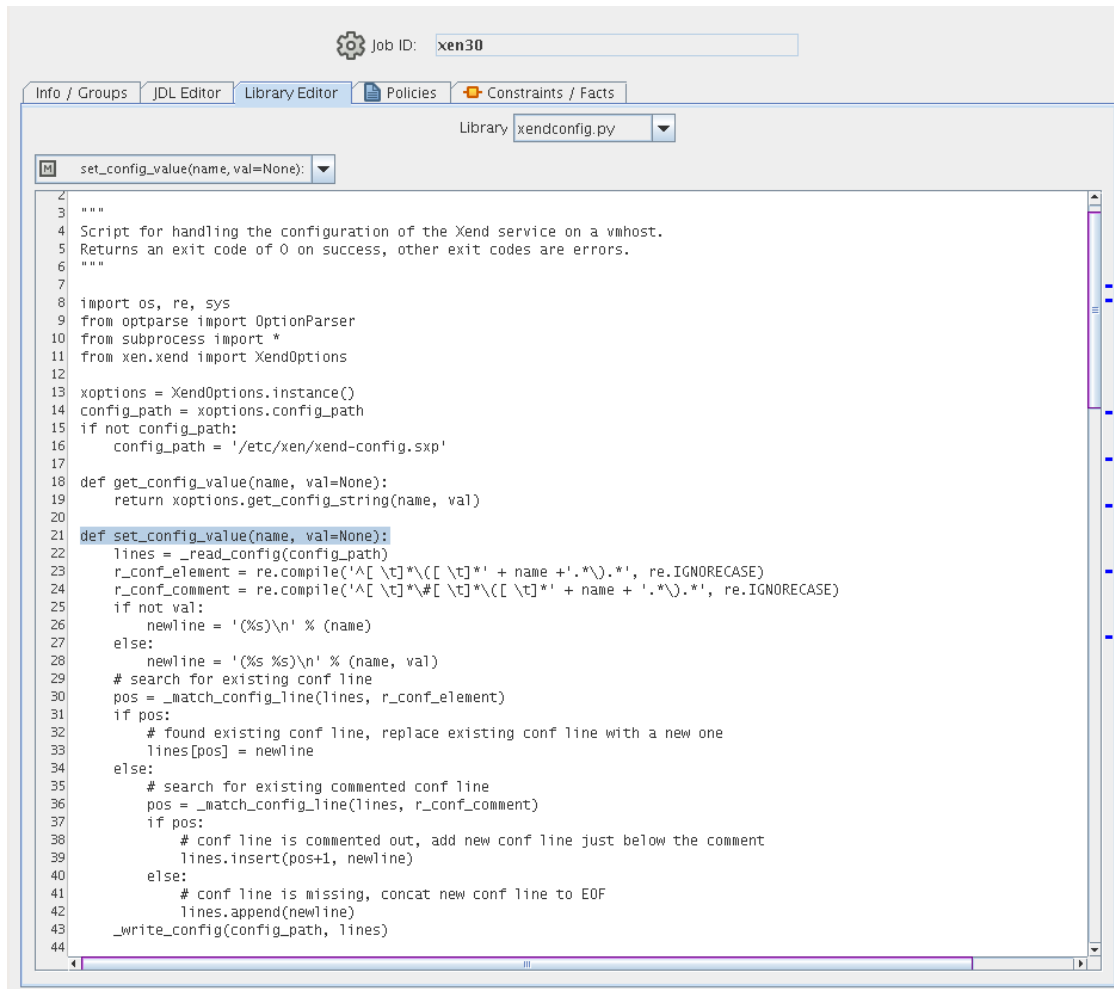


A drop-down list at the top of the editor includes the Java classes and their methods that are bookmarked in the code. Select any of these to go to the location in the code where they are invoked. Clickable colored blocks on the editor scroll bar perform a similar bookmarking function.

4.4 The Job Library Editor Tab

The *Library Editor* tab of the Job admin view opens an editor where you can inspect and modify the different library scripts for a job. The scripts for each job include instructions to the Orchestrate Server for handling job functions.

Figure 4-2 The Job Library Editor



There are two drop-down lists located at the top of the Library Editor view. The first list displays the different libraries for the job, and the second list displays the methods that are bookmarked in the code. Select a method in the second drop-down list to go to the location in the library code where that method is invoked. Clickable colored blocks on the editor scroll bar perform a similar bookmarking function.

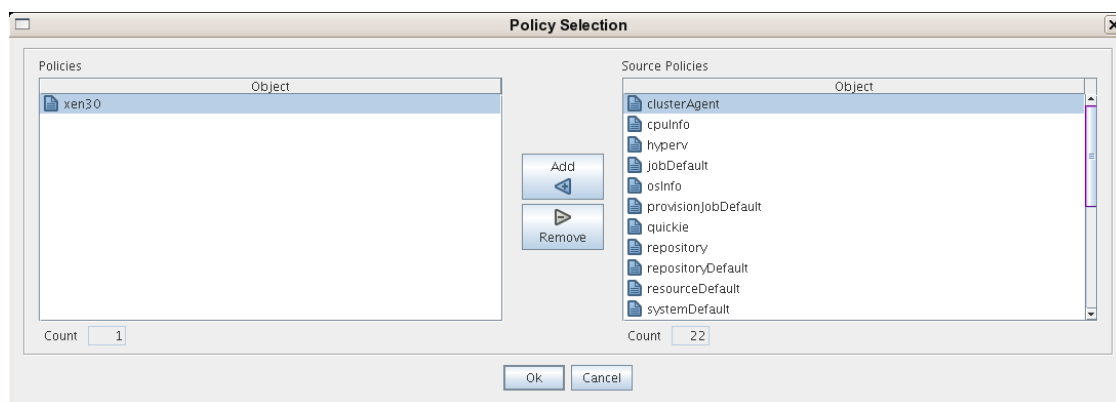
4.5 The Job Policies Tab

The *Policies* tab of the Job admin view opens a page that contains a policy viewer for each of the policies associated with a Job Grid object.

You can modify a policy by using the Policy Grid object. For more information see [Section 12.1, “The Policy Object,”](#) on page 143.

Click *Choose* in the admin view of the Policy viewer to launch a Policy Selection dialog box where you can add or remove individual policies to be applied to the selected Job Grid object.

Figure 4-3 *The Policy Selection Dialog Box*



4.6 The Job Constraints/Facts Tab

The *Constraints/Facts* tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. By changing the policy constraints and fact values for a job, you can change the behavior of the job and how the PlateSpin Orchestrate Server allocates available system resources to it. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator unless they are read-only. Facts with mode r/o have read-only values, which can be viewed (by using the pencil icon), but changes cannot be made.

The Resource Object

A Resource object in the Explorer tree represents a fixed physical machine or a virtual machine (VM) that is managed by PlateSpin Orchestrate. If a resource is running the PlateSpin Orchestrate Agent, that resource can be scheduled for remote execution of a job.

This section includes information about a Resource object that is visible in the Explorer tree and the accompanying Admin view of the Orchestrate Development Client:



- ♦ [Section 5.1, “Resource Groups,” on page 55](#)
- ♦ [Section 5.2, “The Resource Info/Groups Page,” on page 55](#)
- ♦ [Section 5.3, “The Provision Info Page,” on page 82](#)
- ♦ [Section 5.4, “The Resource Log Page,” on page 83](#)
- ♦ [Section 5.5, “The Resource Policies Page,” on page 83](#)
- ♦ [Section 5.6, “The Resource Health Debugger Page,” on page 83](#)
- ♦ [Section 5.7, “The Resource Constraints/Facts Page,” on page 84](#)
- ♦ [Section 5.8, “Resource Object Naming and Renaming,” on page 84](#)

5.1 Resource Groups

Any group object displayed in the Explorer tree represents a collection of similar object types. Groups can also be created automatically, such as the case when a provisioning adapter (PA) discovers a local repository on a VM host. For example, the xen PA, upon discovery of a VM host, automatically creates a local repository for that VM host and places the created repository in a xen repository group. You can also create groups manually in the Development Client, either by clicking the *Actions* menu and choosing *Create Resource Group* or by right-clicking a Resource Group object (anywhere in the Resource hierarchy) and selecting *New Resource Group*.

5.2 The Resource Info/Groups Page

The page that opens under the *Info/Configuration* tab of the Resource admin view includes several collapsible sections on the page where you can configure the general information and attributes of the job.

NOTE: Whenever you make changes to any Grid object, the write icon is superimposed on the object’s icon , signifying that the object has been altered. If you want to save the changes you have made, you need to click the save icon  on the Development Client toolbar.

- ♦ [Section 5.2.1, “The Info Panel,” on page 55](#)
- ♦ [Section 5.2.2, “The Groups Panel,” on page 82](#)

5.2.1 The Info Panel

The following fields on the Information panel provide facts for the Resource object:

- ♦ [“Show Inherited Fact Values Check Box” on page 56](#)

- ♦ [“Resource Information” on page 56](#)
- ♦ [“VM Host Info” on page 60](#)
- ♦ [“Virtual Machine Configuration” on page 61](#)
- ♦ [“Provisioning Information” on page 65](#)
- ♦ [“OS Information” on page 75](#)
- ♦ [“CPU Information” on page 77](#)
- ♦ [“Memory Information” on page 78](#)
- ♦ [“Disk/Network Information” on page 78](#)
- ♦ [“Agent Information” on page 79](#)
- ♦ [“Agent Configuration” on page 80](#)
- ♦ [“Installed Components” on page 82](#)

Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached or inherited policies. These fact values are read only (non-editable).

Resource Information

The Job Control Settings panel on the Info/Groups page includes the following fields:

- ♦ [Resource Type](#)
- ♦ [Resource Enabled](#)
- ♦ [Healthy](#)
- ♦ [Shutting Down](#)
- ♦ [Host Name](#)
- ♦ [Host Fully Qualified Name](#)
- ♦ [Password](#)
- ♦ [Host IP Address](#)
- ♦ [VNC IP Address](#)
- ♦ [VNC Port](#)
- ♦ [Billing Rate](#)
- ♦ [Bill For](#)
- ♦ [Power Factor](#)
- ♦ [Load Average](#)
- ♦ [CPU Load](#)
- ♦ [Joblet Slots](#)
- ♦ [Extra System Joblet Slots](#)
- ♦ [Joblets Active](#)
- ♦ [Became Idle On](#)
- ♦ [Total Joblets Started](#)

- ♦ [Total Completed Joblets](#)
- ♦ [Total Cancelled Joblets](#)
- ♦ [Total Failed Joblets](#)
- ♦ [Total Charge](#)
- ♦ [Total Wall Time](#)
- ♦ [Total Grid Time](#)
- ♦ [Sessions](#)
- ♦ [Provisionable Resource](#)

NOTE: Tool tip text is available when you mouse over any of these fields.

Resource Type: Lets you choose the resource type. If you manually create a resource, you must select the appropriate type.

- ♦ **Fixed Physical:** The node is a physical, hardware-based computer.
- ♦ **VM:** The node is a virtual, software-based container that can run its own operating system and applications as if it were a physical computer.
- ♦ **VM Template:** The node is an image of a server that can be used to create and provision new virtual servers. The template includes a virtual hardware components, a guest operating system, its configuration, and other software applications.

In the Fact Editor, the Resource Type fact is listed as `resource.type`:

```
<fact name="resource.type" value="Fixed Physical" type="String" />
```

Resource Enabled: This check box is selected by default. When it is selected (it has a value of true), the resource is enabled and allowed to accept work.

In the Fact Editor, this fact is listed as `resource.enabled`:

```
<fact name="resource.enabled" value="true" type="Boolean" />
```

Healthy: When this check box is selected (it has a value of true), the resource is considered to be in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy. For more information, see [Appendix A, “Grid Object Health Monitoring,” on page 185](#).

In the Fact Editor, this fact is listed as `resource.health`:

```
<fact name="resource.health" value="true" type="Boolean" />
```

Shutting Down: (Read Only) When this check box is selected (it has a value of true), the node is attempting to shut down, pause, or suspend and does not accept new work.

In the Fact Editor, this fact is listed as `resource.shuttingdown`:

```
<fact name="resource.shuttingdown" value="false" type="Boolean" />
```

Host Name: The network hostname of the resource that is running the Orchestrate Agent. The resource ID and the hostname are often the same, but this is not always the case.

In the Fact Editor, this fact is listed as `resource.hostname`:

```
<fact name="resource.hostname" value="foonode" type="String" />
```

Host Fully Qualified Name: The full network hostname of the resource that is running the Orchestrate Agent.

In the Fact Editor, this fact is listed as `resource.hostname.full`:

```
<fact name="resource.hostname.full" value="foonode.division.company.com" type="String" />
```

Password: The password you want the PlateSpin Orchestrate Agent on this node to use for authentication to the PlateSpin Orchestrate Server.

In the Fact Editor, this fact is listed as `resource.password`.

```
<fact name="resource.password" value="xxx" type="String" />
```

Host IP Address: The network IP address of the resource running the Orchestrate Agent.

In the Fact Editor, this fact is listed as `resource.ip`:

```
<fact name="resource.ip" value="10.255.255.255" type="String" />
```

VNC IP Address: The IP address for a VNC session running on this resource.

In the Fact Editor, this fact is listed as `resource.vnc.ip`:

```
<fact name="resource.vnc.ip" value="" type="String" />
```

VNC Port: The port number for a VNC session running on this resource.

In the Fact Editor, this fact is listed as `resource.vnc.port`:

```
<fact name="resource.vnc.port" value="0" type="Integer" />
```

Billing Rate: The billing rate (in dollars per hour) that you want to charge for this resource running its assigned joblets.

In the Fact Editor, this fact is listed as `resource.billingrate`:

```
<fact name="resource.billingrate" value="1.0000" type="Real" />
```

Bill For: Lets you choose the time scale you want to bill for.

- ♦ **walltime:** The total time for the process to complete.
- ♦ **gcycles:** The normalized average of compute cycles.

In the Fact Editor, this fact is listed as `resource.billfor`:

```
<fact name="resource.billfor" value="walltime" type="String" />
```

Power Factor: (Read Only) The normalized power index of this machine relative to a 2.0 GHz Intel Pentium 4 machine.

In the Fact Editor, this fact is listed as `resource.powerfactor`:

```
<fact name="resource.powerfactor" value="1.0000" type="Real" />
```

Load Average: (Read Only) The load average on this resource as determined with the `uptime` command or other similar methods. The resource is polled every 30 seconds to determine the average.

In the Fact Editor, this fact is listed as `resource.loadaverage`:

```
<fact name="resource.loadaverage" value="0.0000" type="Real" />
```

CPU Load: (Read Only) The percentage of CPU utilization currently used by the resource.

In the Fact Editor, this fact is listed as `resource.cpubload`

```
<fact name="resource.cpubload" value="0" type="Integer" />
```

Joblet Slots: The number of joblets that this resource can run simultaneously.

In the Fact Editor, this fact is listed as `resource.joblets.slots`:

```
<fact name="resource.joblets.slots" value="1" type="Integer" />
```

Extra System Joblet Slots: The number of extra slots you want to be made available to privileged system joblets.

In the Fact Editor, this fact is listed as `resource.joblets.systemslots`:

```
<fact name="resource.joblets.systemslots" value="1" type="Integer" />
```

Joblets Active: (Read Only) The number of joblets that are currently active on this resource.

In the Fact Editor, this fact is listed as `resource.joblets.active`:

```
<fact name="resource.joblets.active" value="0" type="Integer" />
```

Became Idle On: (Read Only) The date and time when the resource became idle. The field displays -1 if the resource is active.

In the Fact Editor, this fact is listed as `resource.becameidle`:

```
<fact name="resource.becameidle" value="7/23/09 5:02 PM" type="Date" />
```

Total Joblets Started: (Read Only) The total number of joblets that have run historically on this resource.

In the Fact Editor, this fact is listed as `resource.history.jobletcount`:

```
<fact name="resource.history.jobletcount" value="8" type="Integer" />
```

Total Completed Joblets: (Read Only) The total number of joblets that have completed historically on this resource.

In the Fact Editor, this fact is listed as `resource.history.jobletcount.completed`:

```
<fact name="resource.history.jobletcount.completed" value="8" type="Integer" />
```

Total Cancelled Joblets: (Read Only) The total number of joblets that have been canceled historically on this resource.

In the Fact Editor, this fact is listed as `resource.history.jobletcount.cancelled`:

```
<fact name="resource.history.jobletcount.cancelled" value="0" type="Integer" />
```

Total Failed Joblets: (Read Only) The total number of joblets that have failed historically on this resource.

In the Fact Editor, this fact is listed as `resource.history.jobletcount.failed`:

```
<fact name="resource.history.jobletcount.failed" value="0" type="Integer" />
```

Total Charge: (Read Only) The cost (in dollars) of all of the joblets run on this resource.

In the Fact Editor, this fact is listed as `resource.history.cost.total`:

```
<fact name="resource.history.cost.total" value="0.0088" type="Real" />
```

Total Wall Time: (Read Only) The total wall time (measured in seconds) that this resource has spent running joblets.

In the Fact Editor, this fact is listed as `resource.history.time.total`:

```
<fact name="resource.history.time.total" value="31" type="Integer" />
```

Total Grid Time: (Read Only) The amount of time (measured in gcycles, which is the normalized average of compute cycles) of all work performed on this resource.

In the Fact Editor, this fact is listed as `resource.history.gcycles.total`:

```
<fact name="resource.history.gcycles.total" value="31" type="Integer" />
```

Sessions: (Read Only) The number of current active sessions (that is, the resource instances with an active agent). The value will be either 1 or 0, unless the object is actually a resource template, in which case it might be greater than 1.

In the Fact Editor, this fact is listed as `resource.sessions`:

```
<fact name="resource.sessions" value="0" type="Integer" />
```

Provisionable Resource: This check box is not selected by default. When it is selected (it has a value of true), you specify that this resource is a provisionable type. Currently, only a VM resource and a VM template resource are provisionable.

In the Fact Editor, this fact is listed as `resource.provisionable`:

```
<fact name="resource.provisionable" value="false" type="Boolean" />
```

VM Host Info


The settings in this section of the Info/Groups page are available when the resource is a VM host.

- ♦ [VM Host Containers](#)
- ♦ [VM Host Repositories](#)

VM Host Containers: A list of VM host containers that are supported by this resource. The list is aggregated from the VM host containers.

In the Fact Editor, this fact is listed as an array:


```
<fact name="resource.vmhosts">
  <array>
    <string>host1slesx_xen30</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor, where you can add or remove VM host containers for the array of element choices.

VM Host Repositories: A list of VM host repositories visible to this resource. The list is aggregated from the VM host repositories.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.repositories">
  <array>
    <string>zos</string>
    <string>vmh3slesx</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor, where you can add or remove VM host repositories to the array of element choices.

Virtual Machine Configuration

The settings in this section of the Info/Groups page are available when the resource is a VM:

- ◆ [Under Construction](#)
- ◆ [VM Vendor](#)
- ◆ [VM UUID](#)
- ◆ [VM Version](#)
- ◆ [Default Storage Repository](#)
- ◆ [Virtual NICs](#)
- ◆ [Virtual NIC Networks](#)
- ◆ [Virtual Disks](#)
- ◆ [Virtual Disk Repositories](#)
- ◆ [Moveable Virtual Disk Repositories](#)
- ◆ [Unmoveable Virtual disk Repositories](#)
- ◆ [Storage Location in Repository](#)
- ◆ [VM Files](#)
- ◆ [Required VM Memory](#)
- ◆ [Required VM Memory Overhead](#)
- ◆ [Required Total VM Memory](#)
- ◆ [Host CPU Architecture](#)
- ◆ [Requires Host HVM Support](#)
- ◆ [Host CPU % Weight](#)
- ◆ [Host CPU Number](#)
- ◆ [Moveable Disk Size](#)
- ◆ [Allow VM Migration](#)
- ◆ [VM Host Ranking](#)
- ◆ [Construction Specification](#)

Under Construction: This check box is not selected by default. When it is selected (it has a value of true), the VM is currently in the process of being created and cannot be provisioned.

In the Fact Editor, this fact is listed as `resource.vm.underconstruction`:

```
<fact name="resource.vm.underconstruction" value="false" type="Boolean" />
```

VM Vendor: The vendor name of the hypervisor that has provided the virtual machine.

In the Fact Editor, this fact is listed as `resource.vm.vendor`:

```
<fact name="resource.vm.vendor" value="xen" type="String" />
```

VM UUID: The vendor and adapter-specific UUID of the resource. You should edit this value only if you are manually creating a Resource object.

In the Fact Editor, this fact is listed as `resource.vm.uuid`:

```
<fact name="resource.vm.uuid" value="237e9975-xxx15-yy1122-7c62-bf6d23d3a049" type="String" />
```

VM Version: This fact is no longer used.

In the Fact Editor, this fact is listed as `resource.vm.version`:

```
<fact name="resource.vm.version" value="0" type="Integer" />
```

Default Storage Repository: Lets you choose the repository where the images of this VM disk and other configuration files are currently stored or where they will be stored.

In the Fact Editor, this fact is listed as `resource.vm.repository`:

```
<fact name="resource.vm.repository" value="vmh1slesx" type="String" />
```

Virtual NICs: The virtual network interface cards (VNICs) that make up this VM. The list is aggregated from the VNIC containers.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.vm.vnics">
  <array>
    <string>win2003_vnic1</string>
  </array>
</fact>
```

Virtual NIC Networks: The networks associated with the VM network interfaces.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.vm.networks">
  <array>
    <string>eth1</string>
  </array>
</fact>
```

Virtual Disks: The list of virtual disks that make up this VM.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.vm.vdisks">
  <array>
    <string>websrvr_vdisk1</string>
  </array>
</fact>
```

Virtual Disk Repositories: The repositories where the VM disk images are stored..

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.vm.repositories">
  <array>
    <string>zos</string>
  </array>
</fact>
```

Moveable Virtual Disk Repositories: The repositories where the moveable VM disk images are stored.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.vm.repositories.moveable">
  <array>
    <string>zos</string>
  </array>
</fact>
```

Unmoveable Virtual Disk Repositories: The repositories where the unmoveable VM disk images are stored.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.vm.repositories.unmoveable">
  <array type="String">
  </array>
</fact>
```

Storage Location in Repository: The file system location (either absolute or relative to repository.location) of the VM file(s).


In the Fact Editor, this fact is listed as resource.vm.basepath:

```
<fact name="resource.vm.basepath" value="vm/websrvr" type="String" />
```

VM Files: The files that make up this VM. The dictionary key (String) represents the file type (adapter specific). The value is the file path either absolute or relative to repository.location of the resource.vm.repository.

In the Fact Editor, this fact is listed as a dictionary:

```
<fact name="resource.vm.files">
  <dictionary>
    <dictelement key="config">
      <string>/var/lib/xen/images/win2kbuild/config.xen</string>
    </dictelement>
  </dictionary>
</fact>
```

You can edit this array by clicking the  button to open an array editor, where you can add or remove ranking specifications for the array of element choices.

Required VM Memory: The amount (measured in MB) of virtual memory required for this VM image.

In the Fact Editor, this fact is listed as `resource.vm.memory`:

```
<fact name="resource.vm.memory" value="1024" type="Integer" />
```

Required VM Memory Overhead: The amount (measured in MB) of virtual memory overhead required for this VM image to provision.

In the Fact Editor, this fact is listed as `resource.vm.memory.overhead`:

```
<fact name="resource.vm.memory.overhead" value="70" type="Integer" />
```

Required Total VM Memory: The total amount (measured in MB) of virtual memory required for this VM image.

In the Fact Editor, this fact is listed as `resource.vm.memory.total`:

```
<fact name="resource.vm.memory" value="402" type="Integer" />
```

Host CPU Architecture: The type of CPU architecture required by this VM. You should edit these values only when you are manually creating a Resource object.

Possible types include:

- ♦ x86
- ♦ x86_64
- ♦ sparc
- ♦ ppc
- ♦ mips
- ♦ alpha

In the Fact Editor, this fact is listed as `resource.vm.cpu.architecture`:

```
<fact name="resource.vm.cpu.architecture" value="x86" type="String" />
```

Requires Host HVM Support: This check box is selected by default. When it is selected (it has a value of true), this VM requires host HVM support. The setting is required when you want to perform paravirtualization; otherwise, only full virtualization is possible.

In the Fact Editor, this fact is listed as `resource.vm.cpu.hvm`:

```
<fact name="resource.vm.cpu.hvm" value="true" type="Boolean" />
```

Host CPU % Weight: The CPU weight (as a percentage of the virtual processor runtime) that you can assign to the virtual processor associated with this VM.

A value of 1.0 represents normal weighting. Setting another VM to a weight of 2.0 means that it would get twice as much CPU runtime as this VM.

In the Fact Editor, this fact is listed as `resource.vm.cpu.weight`:

```
<fact name="resource.vm.cpu.weight" value="1.0000" type="Real" />
```

Host CPU Number: The number of virtual CPUs assigned to this VM.

In the Fact Editor, this fact is listed as `resource.vm.vcpu.number`:

```
<fact name="resource.vm.vcpu.number" value="1" type="Integer" />
```

Moveable Disk Size: The total size (measured in MB) of all the virtual moveable disks for this VM image.

In the Fact Editor, this fact is listed as `resource.vm.vdisksize`:

```
<fact name="resource.vm.vdisksize" value="4096" type="Integer" />
```

Allow VM Migration: This check box is selected by default. When it is not selected (it has a value of false), PlateSpin Orchestrate prevents migration of the VM to another potential VM host.


In the Fact Editor, this fact is listed as `resource.vm.migratable`:

```
<fact name="resource.vm.migratable" value="true" type="Boolean" />
```

VM Host Ranking: This list box includes the ranking specifications used to select suitable VM hosts. The element syntax is `fact/order`, where `order` is either ascending or descending.

In the Fact Editor, this fact is listed as an array:


```
<fact name="resource.vm.vmhost.rankby">
  <array>
    <string>vmhost.vm.placement.score/a</string>
    <string>vmhost.loadindex.slots/a</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor, where you can add or remove ranking specifications for the array of element choices. A trailing `/a` indicates an ascending sort order. A trailing `/d` indicates a descending sort order.

Construction Specification: The *VM Builder Specifications* field in this section displays a list of specifications that were used to build this VM. These specifications are interpreted by the provisioning adapter. You should edit these values by using the PlateSpin Orchestrate VM Client.

In the Fact Editor, this fact is listed as a dictionary:

```
<fact name="resource.vm.spec">
  <dictionary>
    <dictelement key="ssd">
      <string>ddd</string>
    </dictelement>
  </dictionary>
</fact>
```

You can edit the dictionary by clicking the  button to open an attribute editor where you can add or remove dialog box and then expand the map to open the details of the dictionary.

Provisioning Information

The settings on this section of the Info/Groups panel are not available unless the resource you select is a VM.

- ♦ [Provisioning Job](#)

- ◆ [Provisioned Instances](#)
- ◆ [Cloned Instances](#)
- ◆ [Instances](#)
- ◆ [Max Provisioned Instances](#)
- ◆ [Agent Shutdown Timeout](#)
- ◆ [Default Agent Idle Timeout](#)
- ◆ [Host Wait Timeout](#)
- ◆ [Preferred Host Wait](#)
- ◆ [Recommended Host](#)
- ◆ [Debug Provision Log](#)
- ◆ [Parent Template](#)
- ◆ [Current State](#)
- ◆ [Current Host](#)
- ◆ [Current Status](#)
- ◆ [Current Action](#)
- ◆ [Request Time](#)
- ◆ [Start Time](#)
- ◆ [Shutdown Time](#)
- ◆ [Host Wait Time](#)
- ◆ [Managing Job ID](#)
- ◆ [Automatic Provision](#)
- ◆ [Needs Resync](#)

If the Resource object is a VM, it can be automatically personalized for provisioning with information you provide in one of the following subpanels of the *Provisioning Information* panel:

- ◆ [“Linux Autoprep Config:” on page 69](#)
- ◆ [“Windows Sysprep Config” on page 71](#)

Provisioning Job: Lets you select the name of the provisioning job that manages the life cycle of this resource.

In the Fact Editor, this fact is listed as `resource.provisioner.job`:

```
<fact name="resource.provisioner.job" value="xen30" type="String" />
```

Provisioned Instances: The total count of operational instances and provisions in progress.

In the Fact Editor, this fact is listed as `resource.provisioner.count`:

```
<fact name="resource.provisioner.count" value="0" type="Integer" />
```

Cloned Instances: The total count of cloned instances of the template.


In the Fact Editor, this fact is listed as `resource.provisioner.instancecount`:

```
<fact name="resource.provisioner.instancecount" value="0" type="Integer" />
```

Instances: The IDs of the instances of this template resource (if applicable).

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.provisioner.instances">
  <array type="String">
    </array>
  </fact>
```

You can edit this array by clicking the  button to open an array editor, where you can add or remove instance IDs to the array of choices.

Max Provisioned Instances: (For VM templates only) The maximum allowed number of instances of this provisionable resource.

In the Fact Editor, this fact is listed as `resource.provisioner.maxinstances`:

```
<fact name="resource.provisioner.maxinstances" value="1" type="Integer" />
```

Agent Shutdown Timeout: The maximum amount of time, measured in seconds, allowed for this VM to shut down and for the Orchestrate Agent to disconnect from the Orchestrate Server.

In the Fact Editor, this fact is listed as `resource.provisioner.timeout.shutdown`:

```
<fact name="resource.provisioner.timeout.shutdown" value="" type="Integer" />
```

Default Agent Idle Timeout: The maximum amount of time, measured in seconds, allowed for a VM instance to be idle before relaxing reservation policy or shutting down. Behavior depends on mode, and can be overridden by the provision request.

In the Fact Editor, this fact is listed as `resource.provisioner.timeout.idle`:

```
<fact name="resource.provisioner.timeout.idle" value="" type="Integer" />
```

Host Wait Timeout: The maximum amount of time, measured in seconds, to wait for a suitable host before timing out. A value of less than zero (<0), means that the VM waits indefinitely.

In the Fact Editor, this fact is listed as `resource.provisioner.host.maxwait`:

```
<fact name="resource.provisioner.host.maxwait" value="-1" type="Integer" />
```

Preferred Host Wait: The amount of time, measured in seconds, after which some VM Host constraints (for example, whether to move the disk image) are lifted to increase the available pool of hosts. A value of less than zero (<0), means that the VM resource waits indefinitely.

In the Fact Editor, this fact is listed as `resource.provisioner.host.preferredwait`:

```
<<fact name="resource.provisioner.host.preferredwait" value="0" type="Integer" />
```

Recommended Host: The names of VM hosts that you can choose to associate with this VM resource image. You might specify this host when you want a quick VM startup or if you want to change hosts because the original host was suspended. When combined with the `resource.provisioner.host.preferredwait` fact, this fact can lock a VM to one host.

In the Fact Editor, this fact is listed as `resource.provisioner.recommendedhost`:

```
<fact name="resource.provisioner.recommendedhost" value="" type="String" />
```

Debug Provision Log: This check box is not selected by default. When it is selected (it has a value of true), the debug log level in the provisioner is enabled.

In the Fact Editor, this fact is listed as `resource.provisioner.debug`:

```
<fact name="resource.provisioner.debug" value="false" type="Boolean" />
```

Parent Template: The ID of the template resource from which this instance was created. This is only applicable if the template was copied from another template.

In the Fact Editor, this fact is listed as `resource.provision.template`:

```
<fact name="resource.provision.template" value="" type="String" />
```

Current State: The current state of this provisioned instance. The different states include:

- ♦ down
- ♦ suspended
- ♦ up
- ♦ paused
- ♦ unknown (when an administrative action is in process)

In the Fact Editor, this fact is listed as `resource.provision.state`:

```
<fact name="resource.provision.state" value="down" type="String" />
```

Current Host: The ID of the VM host that is currently housing this provisioned resource.

In the Fact Editor, this fact is listed as `resource.provision.vmhost`:

```
<fact name="resource.provision.vmhost" value="vmh6sles_xen30" type="String" />
```

Current Status: (Read Only) The current descriptive status of the provisioned resource.

In the Fact Editor, this fact is listed as `resource.provision.status`:

```
<fact name="resource.provision.status" value="Undefined" type="String" />
```

Current Action: (Read Only) The management action currently in progress on this provisioned resource.

In the Fact Editor, this fact is listed as `resource.provision.currentaction`:

```
<fact name="resource.provision.currentaction" value="" type="String" />
```

Request Time: (Read Only) The time when the last provision (or other administrative action) was requested.

In the Fact Editor, this fact is listed as `resource.provision.time.request`:

```
<fact name="resource.provision.time.request" value="8/24/09 4:36 PM" type="Date" />
```

Start Time: (Read Only) The time when the resource was last successfully provisioned.

In the Fact Editor, this fact is listed as `resource.provision.time.start`:

```
<fact name="resource.provision.time.start" value="12/31/69 4:59 PM" type="Date" />
```

Shutdown Time: (Read Only) The time when the resource was last shut down.

In the Fact Editor, this fact is listed as `resource.provision.time.shutdown`:

```
<fact name="resource.provision.time.shutdown" value="12/31/69 4:59 PM" type="Date" />
```

Host Wait Time: (Read Only) The amount of time (measured in seconds) that this resource has been waiting for or did wait for a suitable host.

In the Fact Editor, this fact is listed as `resource.provision.time.hostwait`:

```
<fact name="resource.provision.time.hostwait" value="0" type="Integer" />
```

Managing Job ID: (Read Only) The current or last Job ID that performed a provisioning action on this resource. This is useful when viewing the job log to monitor specific provisioning actions.

In the Fact Editor, this fact is listed as `resource.provision.jobid`:

```
<fact name="resource.provision.jobid" value="system.xen30.74239" type="String" />
```

Automatic Provision: (Read Only) This check box is not selected by default. When it is selected (it has a value of true), the resource was cloned or provisioned automatically and will be shut down or destroyed automatically.

In the Fact Editor, this fact is listed as `resource.provision.automatic`:

```
<fact name="resource.provision.automatic" value="false" type="Boolean" />
```

Needs Resync: This check box is not selected by default. When it is selected (it has a value of true), you specify that the provisioned resource's state needs to be resynchronized by using the associated provisioning technology at the next opportunity.

In the Fact Editor, this fact is listed as `resource.provision.resync`:

```
<fact name="resource.provision.resync" value="false" type="Boolean" />
```

Linux Autoprep Config:

NOTE: This section displays when a Linux VM is selected.

If any of the fields in this section are blank (that is, undefined), click *Define* to install a fact editor that you can use to define the value.

This section includes the following settings:

- ♦ [Linux Computer Name](#)
- ♦ [Linux Domain](#)

The section also includes a subpanel where you can set [Network Autoprep Configuration](#) information.

Linux Computer Name: This value specifies the host name of a new VM. Enter "*" to indicate that the VM ID is to be used rather than the host name you specify.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.linuxglobal.ComputerName`

```
<fact name="resource.provisioner.autoprep.linuxglobal.ComputerName"
value="afd" type="String" />
```

Linux Domain: This value specifies the domain to which the new VM belongs.

In the Fact Editor, this fact is listed as resource.provisioner.autoprep.linuxglobal.Domain

```
<fact name="resource.provisioner.autoprep.linuxglobal.Domain" value=""
type="String" />
```

Network Autoprep Config


This section includes the following settings:

- ♦ [DNS Server IP Addresses](#)
- ♦ [DNS Suffixes](#)
- ♦ [Gateway IP Addresses](#)

DNS Server IP Addresses: This field displays a list of DNS server IP addresses for name lookup. This is only for cloning/personalize actions.

In the Fact Editor, this fact is listed as an array:


```
<fact name="resource.provisioner.autoprep.DNSServers">
  <array>
    <string>0.0.00.200</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove a server IP address or change its order in the array of element choices.

DNS Suffixes: The list of suffixes to append to a name for lookup. This is only for cloning/personalize actions.

In the Fact Editor, this fact is listed as an array:


```
<fact name="resource.provisioner.autoprep.DNSSuffixes">
  <array>
    <string>afjkd</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove a suffix or change its order in the array of element choices.

Gateway IP Addresses: The list of internet gateways available to this VM. This is only for cloning/personalize actions.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.provisioner.autoprep.Gateways">
  <array>
    <string>afdasadfs</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove the IP address or change its order in the array of element choices.

Windows Sysprep Config

NOTE: This section displays when a Windows VM is selected.

If any of the fields in this section are blank (that is, undefined), click *Define* to install a fact editor that you can use to define the value.

The section includes the following settings/facts:

- ♦ [Change SID](#)
- ♦ [Delete Accounts](#)
- ♦ [Admin Password](#)
- ♦ [Admin Password Plaintext](#)
- ♦ [Timezone](#)
- ♦ [Autologon](#)
- ♦ [Autologon Count](#)
- ♦ [Fullname](#)
- ♦ [Org Name](#)
- ♦ [Computer Name](#)
- ♦ [Product ID](#)
- ♦ [Run Once Command](#)
- ♦ [Workgroup](#)
- ♦ [Domain](#)
- ♦ [Domain Admin](#)
- ♦ [Domain Admin Password](#)
- ♦ [Domain Admin Password Plaintext](#)
- ♦ [Machine Object OU](#)
- ♦ [Machine Password](#)
- ♦ [Machine Password Plaintext](#)
- ♦ [License File Automode](#)
- ♦ [License File Autousers](#)

The section also includes a subpanel where you can set [Network Sysprep Configuration](#) information.

Change SID: The Windows Security ID. If true, sysprep generates a new Security ID.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.options.changeSID`:

```
<fact name="resource.provisioner.autoprep.options.changeSID" value="false"
type="Boolean" />
```

Delete Accounts: If set to true, this fact removes all accounts from the destination VM. If false, it retains existing accounts from the source VM.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.options.deleteAccounts:

```
<fact name="resource.provisioner.autoprep.options.deleteAccounts"
value="true" type="Boolean" />
```

Admin Password: This field displays the IP address for the adapter.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.value:

```
<fact
name="resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.value
" value="klvm" type="String" />
```

Admin Password Plaintext: This field displays the subnet mask for the adapter.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.plainText:

```
<fact
name="resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.plain
Text" value="false" type="Boolean" />
```

Timezone: The time zone of the new VM. See [Microsoft \[GUI Unattended\] \(Sysprep\) product documentation \(http://technet.microsoft.com/en-us/library/cc772783%28WS.10%29.aspx\)](http://technet.microsoft.com/en-us/library/cc772783%28WS.10%29.aspx). (Scroll to the *TimeZone* heading on that page.)

If you do not specify a value for this fact, the default value is 004 (Pacific Standard Time).

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.GuiUnattended.TimeZone:

```
<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.TimeZone"
value="10" type="String" />
```

Autologon: If true, the VM automatically logs into the Administrator account using AdminPassword. If false, logon is prompted.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogon:

```
<fact name="resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogon"
value="true" type="Boolean" />
```

Autologon Count: The limit count for the VM to auto log on with the Administrator account. *AutoLogon* must be True.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogonCount:

```
<fact
name="resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogonCount"
value="2" type="Integer" />
```

Fullname: The user's full name.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.UserData.FullName:


```
<fact name="resource.provisioner.autoprep.sysprep.UserData.FullName"
value="adfk1" type="String" />
```

Org Name: The organization name.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.UserData.OrgName:

```
<fact name="resource.provisioner.autoprep.sysprep.UserData.OrgName"
value="Novell" type="String" />>
```

Computer Name: The VM's new host name. An asterisk (*) means to generate a name based on the source VM name.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.UserData.ComputerName:

```
<fact name="resource.provisioner.autoprep.sysprep.UserData.ComputerName"
value="docdev1" type="String" />
```

Product ID: The Windows product key.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.UserData.ProductID:

```
<fact name="resource.provisioner.autoprep.sysprep.UserData.ProductID"
value="jklaiouga4354" type="String" />
```

Run Once Command: A list of commands that run the first time a user logs on after the new VM is created. Commands are scheduled using the HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\RunOnce registry key.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.GuiRunOnce.Command:

```
<fact name="resource.provisioner.autoprep.sysprep.GuiRunOnce.Command"
value="purge" type="String" />>
```

Workgroup: Windows workgroup name. If joining a domain, use JoinDomain.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.JoinWorkgroup:

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.JoinWorkgroup"
value="prod" type="String" />
```

Domain: Windows domain name. If the VM is joining a workgroup, use JoinWorkgroup. For joining a domain, DomainAdmin and DomainAdminPassword must be defined..

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.JoinDomain:

```
<fact name="resource.provisioner.autoprep.sysprep.Identification.JoinDomain"
value="test" type="String" />>
```

Domain Admin: The Windows domain administrator name.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.DomainAdmin:

```
<fact name="resource.provisioner.autoprep.sysprep.Identification.DomainAdmin"
value="admin" type="String" />
```

Domain Admin Password: The Windows domain administrator account password.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.value:

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.DomainAdminPasswor
d.value" value="cleanwindow" type="String" />
```

Domain Admin Password Plaintext: Select the check box if DomainAdminPassword is in plain text..

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.plainText:

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.DomainAdminPasswor
d.plainText" value="false" type="Boolean" />
```

Machine Object OU: Enter the organizational unit (OU) of the Windows Active Directory machine.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.MachineObjectOU

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.MachineObjectOU"
value="dd" type="String" />
```

Machine Password: Enter the account password for the Windows Active Directory machine.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.MachinePassword.value

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.MachinePassword.va
lue" value="fad" type="String" />
```

Machine Password Plaintext: Select the check box if MachinePassword is in plain text.

In the Fact Editor, this fact is listed as

resource.provisioner.autoprep.sysprep.Identification.MachinePassword.plainText

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.MachinePassword.pl
ainText" value="true" type="Boolean" />
```

License File Automode: Enter either PerServer or PerSeat. If you enter PerServer, AutoUsers must be set.

In the Fact Editor, this fact is listed as

`resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoMode:`

```
<fact
name="resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoMode"
value="PerSeat" type="String" />
```

License File Autousers: The number of client licenses. Use this setting only if AutoMode is PerServer.

In the Fact Editor, this fact is listed as

`resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoUsers:`

```
<fact
name="resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoUsers"
value="33" type="Integer" />
```

Network Sysprep Config


This section includes the following settings:

- ◆ [DNS Suffixes](#)

DNS Suffixes: The list of suffixes to append to a name for lookup. This is only for cloning/ personalize actions.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.provisioner.autoprep.DNSSuffixes">
  <array>
    <string>afjkd</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove a suffix or change its order in the array of element choices.

OS Information

This section includes the following settings:

- ◆ [OS Name](#)
- ◆ [OS Version](#)
- ◆ [OS Version String](#)
- ◆ [OS Architecture](#)
- ◆ [OS Family](#)
- ◆ [OS Type](#)
- ◆ [OS Vendor](#)
- ◆ [OS Vendor Version](#)
- ◆ [OS Vendor String](#)
- ◆ [OS File Path Separator](#)

OS Name: (Read Only) The name of the resource operating system.

In the Fact Editor, this fact is listed as `resource.os.name`:

```
<fact name="resource.os.name" value="Windows" type="String" />
```

OS Version: (Read Only) The version number of the resource operating system. For a VM resource, this fact remains undefined.

In the Fact Editor, this fact is listed as `resource.os.version`:

```
<fact name="resource.os.version.string" value="Microsoft Windows XP [Version 5.1.2600]" type="String" />
```

OS Version String: (Read Only) The operating system vendor full identification string (requires `osInfo` system job).

In the Fact Editor, this fact is listed as `resource.os.version.string`:

```
<fact name="resource.os.vendor.string" value="Microsoft Windows XP [Version 5.1.2600]" type="String" />
```

OS Architecture: The operating system architecture (for example, *x86*, *amd64*, *i386*, or *sparc*).

In the Fact Editor, this fact is listed as `resource.os.arch`:

```
<fact name="resource.os.arch" value="i386" type="String" />
```

OS Family: The operating system family name (for example, *windows*, *linux*, *solaris*, *unix*, *aix*, *mac*) of the resource, if known.

In the Fact Editor, this fact is listed as `resource.os.family`:

```
<fact name="resource.os.family" value="linux" type="String" />
```

OS Type: This drop-down list lets you select the unique string identifier for each OS release, for example, *sles11*.

In the Fact Editor, this fact is listed as `resource.os.type`:

```
<fact name="resource.os.type" value="sles10" type="String" />
```

OS Vendor: (Read Only) The operating system vendor (SuSE for SUSE Linux Enterprise Server or SUSE Linux Enterprise Desktop).

In the Fact Editor, this fact is listed as `resource.os.vendor`:

```
<fact name="resource.os.vendor" value="SuSE" type="String" />
```

OS Vendor Version: (Read Only) This field displays the vendor-defined version for the operating system (for example, *10* for SUSE Linux Enterprise Server 10).

In the Fact Editor, this fact is listed as `resource.os.vendor.version`:

```
<fact name="resource.os.vendor.version" value="10" type="String" />
```

OS Vendor String: (Read Only) This field displays the full identification for the operating system that is supplied by the vendor. The `osinfo` system job must run for this value to be displayed.

In the Fact Editor, this fact is listed as `resource.os.vendor.string`:

```
<fact name="resource.os.vendor.string" value="Welcome to SUSE Linux Enterprise Server 11 (i586) - Kernel (\1)." type="String" />
```

OS File Path Separator: (Read Only) The resource operating system file separator.

In the Fact Editor, this fact is listed as `resource.os.file.separator`:

```
<fact name="resource.os.file.separator" value="/" type="String" />
```

CPU Information

This section includes the following settings:

- ♦ [Number of CPUs](#)
- ♦ [CPU Speed \(Mhz\)](#)
- ♦ [CPU Vendor](#)
- ♦ [CPU Model](#)
- ♦ [CPU Architecture](#)
- ♦ [CPU HVM Support](#)

Number of CPUs: (Read only) The number of CPUs available for this resource to use. For a VM resource, this fact remains undefined.

In the Fact Editor, this fact is listed as `resource.cpu.number`:

```
<fact name="resource.cpu.number" value="2" type="Integer" />
```

CPU Speed (Mhz): (Read only) The processor speed measured in Mhz. The `cpuinfo` job must run for this value to be displayed. For a VM resource, this fact remains undefined.

In the Fact Editor, this fact is listed as `resource.cpu.mhz`:

```
<fact name="resource.cpu.mhz" value="2594" type="Integer" />
```

CPU Vendor: (Read only) The name of the CPU vendor. The `cpuinfo` system job must run for this value to be displayed. For a VM resource, this fact remains undefined.

In the Fact Editor, this fact is listed as `resource.cpu.vendor`:

```
<fact name="resource.cpu.vendor" value="GenuineIntel" type="String" />
```

CPU Model: (Read only) The full vendor model number of the CPU. The `cpuinfo` system job must run for this value to be displayed. For a VM resource, this fact remains undefined.

In the Fact Editor, this fact is listed as `resource.cpu.model`:

```
<fact name="resource.cpu.model" value="Intel(R) Pentium(R) 4 CPU 2.60GHz" type="String" />
```

CPU Architecture: The CPU architecture (for example, *x86*, *x86_64*, *sparc*) of this resource. For a VM resource, this fact remains undefined.

In the Fact Editor, this fact is listed as `resource.cpu.architecture`:

```
<fact name="resource.cpu.architecture" value="x86" type="String" />
```

CPU HVM Support: This field is marked true if the CPU has hardware virtualization support.

In the Fact Editor, this fact is listed as `resource.cpu.hvm`:

```
<fact name="resource.cpu.hvm" value="false" type="Boolean" />
```

Memory Information

This section includes the following settings:

- ♦ [Virtual Memory \(Mb\)](#)
- ♦ [Virtual Available](#)
- ♦ [Physical Memory \(Mb\)](#)
- ♦ [Physical Available](#)
- ♦ [Swap Memory \(Mb\)](#)
- ♦ [Swap Available](#)

Virtual Memory (Mb): (Read only) The total amount of virtual memory (measured in MB) on the resource. The memInfo system job must run for this value to be displayed

In the Fact Editor, this fact is listed as `resource.memory.virtual.total`:

```
<fact name="resource.memory.virtual.total" value="4060" type="Integer" />
```

Virtual Available: (Read only) The amount of available virtual memory (measured in MB) on the resource. The memInfo system job must run for this value to be displayed.

In the Fact Editor, this fact is listed as `resource.memory.virtual.available`:

```
<fact name="resource.memory.virtual.available" value="19951" type="Integer" />
```

Physical Memory (Mb): (Read only) The total amount of physical memory (measured in MB) on the resource. The memInfo system job must run for this value to be displayed

In the Fact Editor, this fact is listed as `resource.memory.physical.total`:

```
<fact name="resource.memory.physical.total" value="3889" type="Integer" />
```

Physical Available: (Read only) The amount of available physical memory (measured in MB) on the resource. The memInfo system job must run for this value to be displayed.

In the Fact Editor, this fact is listed as `resource.memory.physical.available`:

```
<fact name="resource.memory.physical.available" value="3565" type="Integer" />
```

Swap Memory (Mb): (Read only) The total amount of configured swap space (measured in MB) on the resource. The memInfo system job must run for this value to be displayed.

In the Fact Editor, this fact is listed as `resource.memory.swap.total`:

```
<fact name="resource.memory.swap.total" value="16386" type="Integer" />
```

Swap Available: (Read only) The total amount of free swap space (measured in MB) on the resource. The memInfo system job must run for this value to be displayed.

In the Fact Editor, this fact is listed as `resource.memory.swap.available`:

```
<fact name="resource.memory.swap.available" value="16386" type="Integer" />
```

Disk/Network Information

The facts in the *Disk/Network Information* section of the *Info/Groups* page are not currently functional and are not supported.

Agent Information

This section includes the following settings:

- ♦ [Agent Version](#)
- ♦ [Agent Install Dir](#)
- ♦ [Agent Java Version](#)
- ♦ [Agent Java Runtime](#)
- ♦ [Agent Java Vendor](#)
- ♦ [Agent Java Home Dir](#)
- ♦ [Available Agent Memory](#)
- ♦ [Enhanced Exec Available](#)
- ♦ [Clustered Agent](#)

Agent Version: (Read only) The PlateSpin Orchestrate Agent version and build number that is installed on this resource. The string uses the following syntax:

major.minor.point_build

In the Fact Editor, this fact is listed as `resource.agent.version`:

```
<fact name="resource.agent.version" value="2.0.2_70917" type="String" />
```

Agent Install Dir: (Read only) The name of the home directory of the PlateSpin Orchestrate Agent installation files.

In the Fact Editor, this fact is listed as `resource.agent.home`:

```
<fact name="resource.agent.home" value="/opt/novell/zenworks/zos/agent" type="String" />
```

Agent Java Version: (Read only) The version of the Java JVM currently in use by the PlateSpin Orchestrate Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.version`:

```
<fact name="resource.agent.jvm.version" value="1.5.0_17" type="String" />
```

Agent Java Runtime: (Read only) The version of the Java JVM runtime currently in use by the PlateSpin Orchestrate Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.runtime`:

```
<fact name="resource.agent.jvm.runtime" value="1.5.0_17-b04" type="String" />
```

Agent Java Vendor: (Read only) The name of the vendor of the Java JVM currently in use by the PlateSpin Orchestrate Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.vendor`:

```
<fact name="resource.agent.jvm.vendor" value="Sun Microsystems Inc." type="String" />
```

Agent Java Home Dir: (Read only) The path to the home directory of the Java JVM currently in use by the PlateSpin Orchestrate Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.home`:

```
<fact name="resource.agent.jvm.home" value="/opt/novell/zenworks/zos/agent/jre" type="String" />
```

Available Agent Memory: (Read only) The amount of memory (measured in Mb) available to the PlateSpin Orchestrate Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.memory`:

```
<fact name="resource.agent.jvm.memory" value="127" type="Integer" />
```

Enhanced Exec Available: This check box is selected by default. When it is selected (it has a value of true), it indicates that the PlateSpin Orchestrate Agent installed on this resource is able to use enhanced exec features (as opposed to unsupported agent installs, such as AIX).

In the Fact Editor, this fact is listed as `resource.agent.exec.installed`:

```
<fact name="resource.agent.exec.installed" value="true" type="Boolean" />
```

Clustered Agent: This check box is not selected by default. When you select it (it has a value of true), you specify that the agent is “clustered” on this VM resource. This means that it converts duplicate logins to failover logins.

In the Fact Editor, this fact is listed as `resource.agent.clustered`:

```
<fact name="resource.agent.clustered" value="false" type="Boolean" />
```

Agent Configuration

This section includes the following settings:

- ♦ [Gmond Port](#)
- ♦ [Datagrid Cache TTL](#)
- ♦ [Datagrid Cleanup Interval](#)
- ♦ [Exec Daemon Timeout](#)
- ♦ [Exec As Agent User Only](#)
- ♦ [Cleanup After Joblets](#)
- ♦ [Use Enhanced Exec](#)
- ♦ [Log Level](#)
- ♦ [Debug Logging](#)

Gmond Port: The port that the agent uses for gmond. Port 8649 is the default port. A setting of zero (0) or less means that the value is not read.

In the Fact Editor, this fact is listed as `resource.agent.config.gmond.port`:

```
<fact name="resource.agent.config.gmond.port" value="8649" type="Integer" />
```

Datagrid Cache TTL: The amount of time (measured in minutes) that inactive files should remain in the agent’s datagrid cache. A setting of zero (0) turns off the cache.

In the Fact Editor, this fact is listed as `jresource.agent.config.datagrid.cache.lifetime`:


```
<fact name="resource.agent.config.datagrid.cache.lifetime" value="1440"
type="Integer" />
```

Datagrid Cleanup Interval: The amount of time (measured in minutes) that the Orchestrate Server should wait between cleanup sweeps of the agent's datagrid cache.

In the Fact Editor, this fact is listed as

`resource.agent.config.datagrid.cache.cleanupinterval:`

```
<fact name="resource.agent.config.datagrid.cache.cleanupinterval" value="60"
type="Integer" />
```

Exec Daemon Timeout: The amount of time (measured in seconds) that the enhanced exec daemon is to remain running. A setting of zero (0) specifies that the daemon is to remain running. The exec daemon is the non-Java component of the agent that is responsible for executing commands remotely.

In the Fact Editor, this fact is listed as `resource.agent.config.exec.daemon.timeout:`

```
<fact name="resource.agent.config.exec.daemon.timeout" value="300"
type="Integer" />
```

Exec As Agent User Only: This check box is selected by default. When you select it (it has a value of true), you specify that the agent is to always run executables as the Agent User only. Selecting this check box overrides any job fact settings (for the `job.joblet.runtime` fact).

In the Fact Editor, this fact is listed as `resource.agent.config.exec.asagentuseronly:`

```
<fact name="resource.agent.config.exec.asagentuseronly" value="true"
type="Boolean" />
```

Cleanup After Joblets: This check box is not selected by default. When you select it (it has a value of true), you specify that the agent on this resource is to clean up temporary directories created for each joblet. You can deselect this check box for debugging purposes; when you select it again, the cleanup process starts again, deleting temporary directories that were created while the setting was deactivated.

In the Fact Editor, this fact is listed as `resource.agent.config.joblet.cleanup:`

```
<fact name="resource.agent.config.joblet.cleanup" value="true" type="Boolean"
/>
```

Use Enhanced Exec: This check box is not selected by default. When you select it (it has a value of true), you specify that the agent on this resource is to use the enhanced exec feature of the agent, which is available for supported agent installations. Marking this fact as false causes the enhanced exec feature not to be used.

In the Fact Editor, this fact is listed as `resource.agent.config.exec.enhancedused:`

```
<fact name="resource.agent.config.exec.enhancedused" value="true"
type="Boolean" />
```

Log Level: Lets you choose the level of agent logging in terms of the amount of detail (that is, the “verbosity”) you want to include in the agent log. The choices include:

- ♦ quiet
- ♦ normal
- ♦ verbose

In the Fact Editor, this fact is listed as `resource.agent.config.loglevel`:

```
<fact name="resource.agent.config.loglevel" value="normal" type="String" />
```

Debug Logging: This check box is not selected by default. When you select it (it has a value of `true`), you activate the debug function in the agent log, which is additive to the log level.

In the Fact Editor, this fact is listed as `resource.agent.config.logdebug`:


```
<fact name="resource.agent.config.logdebug" value="false" type="Boolean" />
```

Installed Components

Applications: A list of the names of applications (including the full version name) that are installed on this resource. This is useful for constraining joblets to run only on a resource with a particular application installed.

In the Fact Editor, this fact is listed as an array:

```
<fact name="resource.installed.apps">
  <array>
    <string>man-pages-2.39-0.9</string>
    <string>xorg-x11-fonts-scalable-6.9.0-50.45</string>
    <string>cifs-mount-3.0.24-2.23</string>
    <string>gdbm-1.8.3-243.2</string>
    <string>libaio-0.3.104-14.2</string>
    <string>libnl-1.0-18.4</string>
    ...
  </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove the application name or change its order in the array of element choices.

5.2.2 The Groups Panel

This section of the Info/Groups page lists the groups of Resource objects in the grid. Click *Choose* to open the Resource Group Selection dialog box. In this dialog box, you can choose which Resource Groups to display in the Explorer Panel by selecting a group and then clicking *Add* or *Remove* to move it to or from the Source Resource Groups list.

5.3 The Provision Info Page

The *Provision Info* tab is displayed only for VM resource objects selected in the Explorer tree. The read-only fields displayed at the top of the Provision Info page summarize information related to the provisioning of the resource, whether that resource is a VM or VM template.

The page has several subtabs that open other pages that display further information about the VM:

- ♦ **Show Log:** A log that includes historical details about the history of the provisioning of the VM.
- ♦ **Host Assignment Log:** A log of VM host assignment constraint errors of the last migration or provision action.

- ♦ **Autoprep Data:** A list of autoprep/sysprep fact overrides used to prepare this VM instance. For more information about autoprep, see “[Understanding Autoprep](#)” in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.
- ♦ **Policy Debugger:** Information to help you debug the policy (and its resulting constraints and facts) associated with this VM. For more information, see [Chapter 14, “The Policy Debugger,”](#) on page 175.
- ♦ **Action History:** A log that includes historical details about the history of the provisioning of the VM. For more information, see “[Provisioning Actions and History](#)” in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

5.4 The Resource Log Page

Open the Resource Log tab to view the contents of the log file for this resource. You can click *Refresh* to update the content. You can also select *Debug Logging* to activate the debug feature as part of the logging or change the *Log Level* to the level of detail that you want.

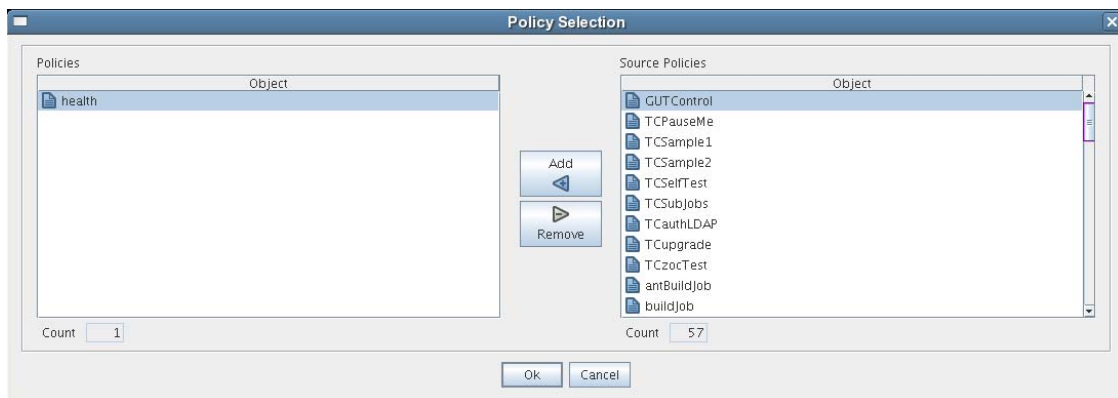
5.5 The Resource Policies Page

The *Policies* tab of the Resource admin view opens a page that contains a policy viewer for each of the policies associated with a Resource Grid object.

You can modify a policy using the Policy Grid object. For more information see [Section 12.1, “The Policy Object,”](#) on page 143.

Click *Choose* in the admin view of the Policy viewer to launch a Policy Selection dialog box where you can add or remove individual policies to be applied to the selected Resource Grid object.

Figure 5-1 The Policy Selection Dialog Box



5.6 The Resource Health Debugger Page

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Chapter A.3, “The Health Debugger,”](#) on page 187.

5.7 The Resource Constraints/Facts Page

The *Constraints/Facts* tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. By changing the policy constraints and fact values for a job, you can change the behavior of the job and how the PlateSpin Orchestrate Server allocates available system resources to it. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator unless they are read-only. Facts with mode *r/o* have read-only values, which can be viewed by using the edit pencil icon, but changes cannot be made.

For information about renaming Resource objects using the Fact Editor on this page, see [Section 5.8, “Resource Object Naming and Renaming,” on page 84](#).

5.8 Resource Object Naming and Renaming

The PlateSpin Orchestrate Resource Grid object type can include resources of various types, including physical machines, virtual machines (VMs), and VM templates, all of which are modeled differently in the Development Client because of their varying roles in the Orchestrate system. Some resource names are generated by the Orchestrate system and can therefore receive generic, arbitrary names such as *mysql-1*, *mysql-2*, and so on. Resources you name at installation time or creation time might also change in their purpose or facilities.

As the quantity of these Resource objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object’s “display name” is visible in the Development Client interface, the VM Client interface, the Server Portal, and in optional *zos* and *zosadmin* commands.

NOTE: Resource object groups (that is, the folders that contain these Resource objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A Resource object’s name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a Resource object in the Orchestrate Development Client using one of three methods:

- ♦ Right-click the Resource object in the Explorer tree, then select *Rename* to allow editing of the display name.
- ♦ Triple-click the Resource object in the Explorer tree to allow editing of the display name.
- ♦ In the *Constraints/Facts* page, where you can select the resource object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you begin to rename using one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the Resource object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the *zos* or *zosadmin* command line, see the [PlateSpin Orchestrate 2.5 Command Line Reference](#).

The VM Host Object

6

A VM host represents a VM host technology or hypervisor (for example, Xen, Hyper-V, and so on) either installed on a physical resource or accessed by it (in the case of VMware). VM host objects can be used when making provisioning decisions for a resource.

PlateSpin Orchestrate also supports the discovery of VMware vSphere clusters used for high availability (HA) in a VMware environment or managed by the VMware Distributed Resource Scheduler (DRS) (after an Orchestrate Agent has been deployed into such an environment). In this scenario, Orchestrate also allows you to determine when actions have taken place outside of Orchestrate, such as when DRS moves a VM to an alternate host in the cluster or when an administrator moves a VM into a different resource pool (see [“Setting Up Orchestrate to Accommodate VMware DRS Clustering and Updates”](#) in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*).

Although the VM host and the VM host Cluster are regarded as two different types of VM host object, and have differing icons, the discovered clusters are represented in the Explorer tree of the Development Client as VM host objects.

NOTE: The Development Client interface (that is, the fields in the admin view) for a VM host and a VM host Cluster are nearly identical. Facts unique to the VM host Cluster are listed in [Section 6.7, “Unique VM Host Cluster Facts,”](#) on page 91.



This section includes the following information:

- ♦ [Section 6.1, “The Info Page,”](#) on page 85
- ♦ [Section 6.2, “The Policies Page,”](#) on page 90
- ♦ [Section 6.3, “The Health Debugger Page,”](#) on page 90
- ♦ [Section 6.4, “The Constraints/Facts Page,”](#) on page 90
- ♦ [Section 6.5, “The Action History Page,”](#) on page 90
- ♦ [Section 6.6, “VM Host Object Naming and Renaming,”](#) on page 91
- ♦ [Section 6.7, “Unique VM Host Cluster Facts,”](#) on page 91
- ♦ [Section 6.8, “vCPU Slots for VM Hosts,”](#) on page 94

6.1 The Info Page

The page that opens under the *Info* tab includes several collapsible sections on the page where you can configure the general information and attributes of the VM host.

- ♦ [Section 6.1.1, “Show Inherited Fact Values Check Box,”](#) on page 86
- ♦ [Section 6.1.2, “VM Host Information Panel,”](#) on page 86
- ♦ [Section 6.1.3, “Provisioning Adapter Config Panel,”](#) on page 88
- ♦ [Section 6.1.4, “Guest VM Monitor Information Panel,”](#) on page 89

NOTE: Whenever you make changes to any Grid object, the write icon  is superimposed on the object's icon, signifying that the object has been altered. If you want to save the changes you have made, you need to click the save icon  on the Development Client toolbar.

6.1.1 Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached or inherited policies. These fact values are read only (non-editable), although you can use the Policy Editor to modify the policy values themselves if you want to.

6.1.2 VM Host Information Panel

The *VM Host Information* panel on the Info page includes the following fields:

NOTE: Tool tip text is available when you mouse over any of these fields.

Physical Resource: (Read Only) The name of the resource that houses this VM host container.

In the Fact Editor, this fact is listed as `vmhost.resource`:

```
<fact name="vmhost.resource" value="vmh7sles" type="String" />
```

VM Host Type: This field displays a read-only fact, as discovered by PlateSpin Orchestrate. It identifies the VM host as a regular VM host (*vmhost*) or as a VMware cluster (*vmhostcluster*).

In the Fact Editor, this fact is listed as `vmhost.type`:

```
<fact name="vmhost.type" value="vmhostcluster" type="String" />
```

VmHost Cluster: (Conditional) This field displays a read-only fact, as discovered by PlateSpin Orchestrate. It identifies the VM host Cluster to which this VM host belongs. The field is displayed on the information page only when the VM host is a member of a cluster.

In the fact editor, this fact is listed as `vmhost.cluster`:

```
<fact name="vmhost.cluster" value="esx35_cluster_vsphere" type="String" />
```

Enabled: This check box is selected by default. When it is selected (it has a value of true), the VM host is enabled, which means that VM instances can be provisioned on it.

In the Fact Editor, this fact is listed as `vmhost.enabled`:

```
<fact name="vmhost.enabled" value="true" type="Boolean" />
```

Online: When this check box is selected (it has a value of true), the agent on the physical resource is online.

In the Fact Editor, this fact is listed as `vmhost.online`:

```
<fact name="vmhost.online" value="true" type="Boolean" />
```

Healthy: This check box is selected by default. When it is selected (it has a value of true), the VM host is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy. For more information, see [Appendix A, “Grid Object Health Monitoring,” on page 185](#)

In the Fact Editor, this fact is listed as `vmhost.health`:

```
<fact name="vmhost.health" value="false" type="Boolean" />
```

Shutting Down: When this check box is selected (it has a value of true), the VM host is attempting to shut down and does not accept provisioning requests.

In the Fact Editor, this fact is listed as `vmhost.shuttingdown`:

```
<fact name="vmhost.shuttingdown" value="false" type="Boolean" />
```

Location: For the vsphere provisioning adapter, this is the `ManagedObjectReference` path for this VM host. For other provisioning adapters, this is an optional description of the physical location of the VM host.

In the Fact Editor, this fact is listed as `vmhost.location`:

```
<fact name="vmhost.location" value="" type="String" />
```

Supports VM Migration: When this check box is selected (it has a value of true), the VM host can support VM migration. The state of this fact can also depend on the migration capabilities of the provisioning adapter used to provision the VM.

In the Fact Editor, this fact is listed as `vmhost.migration`:

```
<fact name="vmhost.migration" value="true" type="Boolean" />
```

Supports H/W HVM: When this check box is selected (it has a value of true), the hypervisor on the VM host can support hardware virtualization.

In the Fact Editor, this fact is listed as `vmhost.hvm`:

```
<fact name="vmhost.hvm" value="false" type="Boolean" />
```

Accounting Group: The default VM host group that you want to be adjusted for VM tracking statistics.

In the Fact Editor, this fact is listed as `vmhost.accountinggroup`:

```
<fact name="vmhost.accountinggroup" value="all" type="String" />
```

Max Hosted VMs: The maximum number of VM instances allowed on this VM host.

In the Fact Editor, this fact is listed as `vmhost.maxvmslots`:

```
<fact name="vmhost.maxvmslots" value="8" type="Integer" />
```

Max Hosted vCPUs: The maximum number of virtual CPUs that this VM host can support.

In the Fact Editor, this fact is listed as `vmhost.vcpu.max`:

```
<fact name="vmhost.vcpu.max" value="8" type="Integer" />
```

For more information about vCPU slots, see [Section 6.8, “vCPU Slots for VM Hosts,” on page 94](#).

Max Virtual Memory: The amount of memory (measured in MB) available to hosted VMs.


In the Fact Editor, this fact is listed as `vmhost.memory.max`:

```
<fact name="vmhost.memory.max" value="1000" type="Integer" />
```

Repositories: The list of repositories (VM disk stores) that are visible to this VM host.

In the Fact Editor, this fact is listed as an array:


```
<fact name="vmhost.repositories">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the  button to open the Choose Grid Objects dialog box, where you can add, remove, or edit repositories in an array of repository choices.

Available VM Resource Groups: This field displays a list of resource groups containing VMs that are allowed to run on this VM host.

In the Fact Editor, this fact is listed as an array:

```
<fact name="vmhost.vm.available.groups">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the  button to open the Choose Grid Objects dialog box, where you can add, remove, or edit the resource groups (element values) in an array of choices.

Managing Job: The ID of a running job that manages VM operations on this VM host. When this field is completed, the VM Manager prevents other jobs from initiating provisioning actions. The fact is cleared when the managing job ends.

In the Fact Editor, this fact is listed as `vmhost.controllingjob`:

```
<fact name="vmhost.controllingjob" value="" type="String" />
```

Needs Resync: When this check box is selected (it has a value of true), you specify that, at the next opportunity, this VM host is to be probed to resynchronize all the VMs that are managed here.

In the Fact Editor, this fact is listed as `vmhost.resync`:

```
<fact name="vmhost.resync" value="false" type="Boolean" />
```

6.1.3 Provisioning Adapter Config Panel

Adapter Job Name: The name of the provisioning adapter job that manages VM discovery on this host. Do not change this value unless you have implemented your own discovery job.

In the Fact Editor, this fact is listed as `vmhost.provisioner.job`:

```
<fact name="vmhost.provisioner.job" value="vsphere" type="String" />
```

Username: (Optional) The username required for provisioning on the VM host.

In the Fact Editor, this fact is listed as `vmhost.provisioner.username`:


```
<fact name="vmhost.provisioner.username" value="" type="String" />
```

Password: (Optional) The password required for provisioning on the VM host.

In the Fact Editor, this fact is listed as `vmhost.provisioner.password`:

```
<fact name="vmhost.provisioner.password" value="" type="String" />
```

6.1.4 Guest VM Monitor Information Panel

Current VM Count: (Read Only) The current number of active VM instances.

In the Fact Editor, this fact is listed as `vmhost.vm.count`:

```
<fact name="vmhost.vm.count" value="0" type="Integer" />
```

Available vCPUs: The number of vCPUs available on this VM host.

In the Fact Editor, this fact is listed as `vmhost.vcpu.available`:

```
<fact name="vmhost.vcpu.available" value="8" type="Integer" />
```

For more information, see [Section 6.8, “vCPU Slots for VM Hosts,” on page 94](#).

Available Virtual Memory: The amount of memory (measured in MB) available to new VMs.


In the Fact Editor, this fact is listed as `vmhost.memory.available`:

```
<fact name="vmhost.memory.available" value="1000" type="Integer" />
```

VM Image Counts: The dictionary of running instance counts for each running VM template.

In the Fact Editor, this fact is listed as a dictionary:

```
<fact name="vmhost.vm.templatecounts">
  <dictionary>
    <dictelement key="ads">
      <time>12:00 AM</time>
    </dictelement>
  </dictionary>
</fact>
```

You can edit the dictionary elements by clicking the  button to open the VM Image Counts dialog box, then adding or removing the names in the dictionary.

Running VM Instances: (Read Only) A list of active VM instances.

In the Fact Editor, this fact is listed as an array:

```
<fact name="vmhost.vm.instanceids">
  <array type="String">
  </array>
</fact>
```

Load Index (Slots): (Read Only) The current loading index of resource slots, which is a ratio of the active hosted VMs to the specified maximum number of VMs allowed on this host. Each provision VM takes up one slot. For more information, see [Max Hosted VMs](#).

In the Fact Editor, this fact is listed as `vmhost.loadindex.slots`:

```
<fact name="vmhost.loadindex.slots" value="0.1250" type="Real" />
```

Load Index (Memory): (Read Only) The current loading index for memory, which is a ratio of the virtual memory consumed on this VM host to the specified maximum amount of memory allocated to this host.

In the Fact Editor, this fact is listed as `vmhost.loadindex.virtualmemory`:

```
<fact name="vmhost.loadindex.virtualmemory" value="0.0000" type="Real" />
```

6.2 The Policies Page

The *Policies* tab opens a page that contains a policy viewer for each of the policies associated with a Grid object.

NOTE: You can edit a policy by right-clicking a policy icon and selecting *Edit Policy*. Remember to click the save icon when your changes are complete..

6.3 The Health Debugger Page

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Chapter A.3, “The Health Debugger,” on page 187](#).

6.4 The Constraints/Facts Page

The Constraints/Facts tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. By building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any system resources by managing an object’s facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode `r/o` have read-only values, which can be viewed (using the pencil icon) but changes cannot be made.

For information about using the Fact Editor on this page to rename VM host objects, see [Section 6.6, “VM Host Object Naming and Renaming,” on page 91](#)

6.5 The Action History Page

The *Action History* tab is displayed in the administrative view of the Repository object. When you select the *Action History* tab, a table displays on the Action History page with a list of the history for all VM provisioning actions performed on this Grid object.

The Orchestrate Server must be connected to an audit database for the *Include Audit Database* check box to be available. If the *Include Audit Database* check box is selected in this view, the action status is not polled. Click the refresh icon in the toolbar to retrieve and display fresh data.

For more details about the information listed on the Action History page, see “[Action History in Admin Views of the Development Client](#)” in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

6.6 VM Host Object Naming and Renaming

Some VM host names (or VM host Cluster names) are generated by the Orchestrate system and can therefore receive generic, arbitrary names such as `host2_demoAdapter`, `host3_demoAdapter`, and so on. VM hosts you name at creation time might also change later in their purpose or facilities.

As the quantity of these VM host objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object's "display name" is visible in the Development Client interface, the VM Client interface, the Server Portal, and in optional `zos` and `zosadmin` commands.

NOTE: Resource object groups (that is, the folders that contain these VM host objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A VM host object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a VM host object (or a VM host Cluster object) in the Orchestrate Development Client using one of three methods:

- ♦ Right-click the VM host object in the Explorer tree, then select **Rename** to allow editing of the display name.
- ♦ Triple-click the VM host object in the Explorer tree to allow editing of the display name.
- ♦ In the *Constraints/Facts* page, where you can select the VM host object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you begin to rename using one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the VM host object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the `zos` or `zosadmin` command line, see the [PlateSpin Orchestrate 2.5 Command Line Reference](#).

6.7 Unique VM Host Cluster Facts

There are several VM host Cluster related facts that are not found in a regular VM host object. This section contains detail about those facts.

- ♦ [Section 6.7.1, "Orchestrate Facts in the VM Host Cluster Object," on page 91](#)
- ♦ [Section 6.7.2, "Orchestrate Facts in a VM Host Residing in a Cluster," on page 92](#)
- ♦ [Section 6.7.3, "Orchestrate Facts in VMs Hosted in Clusters," on page 93](#)

6.7.1 Orchestrate Facts in the VM Host Cluster Object

Any vSphere clusters discovered by PlateSpin Orchestrate are listed in the Development Client as members of a convenience group (for example, a group named `clusters_vsphere`). The following table lists the read-only, cluster-related facts in a VM host Cluster object.

Table 6-1 Cluster-Related Facts in a Cluster Object

Fact Name	Type	Description
<code>vmhost.cluster.vmhosts</code>	String[]	This string array lists all of the VM hosts that are members of this cluster.
<code>vmhost.location</code>	String	The data center Managed Object Reference (MOR) path to the cluster.
<code>vmhost.vphere.cluster.das.admission_control_enabled</code>	Boolean	<p>If this fact value is true, the VM host cluster has HA configuration to use admission control.</p> <hr/> <p>NOTE: HA was originally named “Dynamic Availability Service” (DAS) in VMware. This is the origin of the <code>.das</code> prefix.</p>
<code>vmhost.vphere.cluster.das.enabled</code>	Boolean	If this fact value is true, HA is turned on in this cluster.
<code>vmhost.vphere.cluster.drs.allow_behavior_override</code>	Boolean	Whether VMs can specify their own placement behavior.
<code>vmhost.vphere.cluster.drs.default_behavior</code>	String	Specifies manual, partially automated, or fully automated VM placement.
<code>vmhost.vphere.cluster.drs.enabled</code>	Boolean	If this fact value is true, DRS is turned on in this cluster.
<code>vmhost.vphere.pools</code>	Dictionary	The mapping of the pool MOR path to dictionaries of pool configuration values.
<code>vmhost.type</code>	String	The “type” differentiation for clusters vs. VM hosts (the value is either <code>vmhost</code> or <code>vmhostcluster</code>).

6.7.2 Orchestrate Facts in a VM Host Residing in a Cluster

The following table lists the read-only, cluster-related facts in a virtual or physical machine with an installed Orchestrate Agent and residing in a cluster.

Table 6-2 Cluster-Related Facts in a VM Host Residing in a vSphere Cluster

Fact Name	Type	Description
<code>vmhost.cluster</code>	String[]	The ID of the cluster that contains this host.
<code>vmhost.type</code>	String	The “type” differentiation for clusters vs. VM hosts (the value is either <code>vmhost</code> or <code>vmhostcluster</code>).

6.7.3 Orchestrate Facts in VMs Hosted in Clusters

The following table lists the read-only, cluster-related facts in a VM hosted in vSphere clusters.

Table 6-3 *Cluster-Related Facts in VMs Hosted in Clusters*

Fact Name	Type	Description
<code>resource.provision.vmhost</code>	String	The ID of the VM host cluster or VM host that contains this VM. When provisioning or migrating to a cluster, this value is initially the cluster ID. On subsequent resync or discovery or when using the VsphereUpdate daemon, this is set to the value of the selected VM host within the cluster.
<code>resource.vm.pool</code>	String	The MOR path of the resource pool that contains this VM. This value is also a key in the <code>vmhost.vsphere.pools</code> dictionary.
<code>resource.vm.vmhost.location</code>	String	The MOR path of the VM host cluster that contains this VM. This is always the cluster if the VM host is in a cluster; otherwise, it is the MOR path of the VM host itself.
<code>resource.vm.vsphere.cpu.limit</code>	Integer	The maximum amount (in MHz) of CPU resources to be used by this VM.
<code>resource.vm.vsphere.cpu.reservation</code>	Integer	The minimum amount (in MHz) of CPU resources guaranteed to this VM.
<code>resource.vm.vsphere.cpu.shares.level</code>	String	The relative amount (assigned a value of low, normal, high, or custom) of CPU allocated for this VM.
<code>resource.vm.vsphere.cpu.shares.custom</code>	Integer	The custom amount of relative CPU for this VM. This fact is valid only when the value for the <code>resource.vm.vsphere.cpu.shares.level</code> fact is "custom."
<code>resource.vm.vsphere.memory.limit</code>	Integer	The maximum amount (measured in MB) of memory resources to be used by this VM.
<code>resource.vm.vsphere.memory.reservation</code>	Integer	The minimum amount (measured in MB) of memory resources guaranteed to this VM.
<code>resource.vm.vsphere.memory.shares.level</code>	String	The relative amount (assigned a value of low, normal, high, or custom) of memory specified for this VM.

Fact Name	Type	Description
<code>resource.vm.vsphere.memory.share</code> <code>s.custom</code>	Integer	The custom amount of relative memory for this VM. This fact is valid only when the value for the <code>resource.vm.vsphere.memory.shares.level</code> fact is "custom."

6.8 vCPU Slots for VM Hosts

In PlateSpin Orchestrate, a vCPU represents a logical CPU. It provides a way to set up limits for allocating CPUs on VM hosts. These limits let you specify how many vCPUs should be hosted by each VM host so that you can control how much CPU processing power is available. If all vCPUs are in use, a subsequent provision can be denied or made to wait. This lets you ensure the quality of service you want to maintain in the data center.

When a VM is provisioned, Orchestrate runs a constraint check on every suitable VM host to determine if the number of available vCPUs on the VM host is sufficient for a VM. If a VM host with sufficient available vCPUs is not available, the provision request waits until one becomes available or (depending on VM facts) the request is denied.

Using vCPU facts differs from using the existing slot fact (`vmhost.maxvmslots`). The `maxvmslots` fact provides basic control of the number of VMs allocated to a VM host — useful for limiting VMs because of license restrictions or for generally limiting the VMs being managed. The vCPU facts are similar to the memory limit facts, giving you more control to avoid overloading a VM host and letting you ensure quality of service.

This section includes information about how the vCPU facts are used in PlateSpin Orchestrate.

- ♦ [Section 6.8.1, “Configuring vCPUs on VM Hosts,” on page 94](#)
- ♦ [Section 6.8.2, “Configuring vCPUs on VM Host Clusters,” on page 95](#)
- ♦ [Section 6.8.3, “Configuring vCPUs on VMs,” on page 95](#)

6.8.1 Configuring vCPUs on VM Hosts

There are two vCPU facts displayed on the VM host Info page in the Orchestrate Development Client:

- ♦ **Max Hosted vCPUs:** This value (Integer) represents the maximum number of vCPUs that the VM host can support. The fact name is `vmhost.vcpu.max`.
- ♦ **Available vCPUs:** This value (Integer) represents the number of virtual CPUs available on this host. The fact name is `vmhost.vcpu.available`.

The `vmhost.vcpu.available` value changes when a VM is provisioned or shut down on that VM host. If the `vmhost.vcpu.max` fact is set to -1 (unlimited), the `vmhost.vcpu.available` value changes to -1 (unlimited). When set to unlimited, no counting occurs, so PlateSpin Orchestrate does not check vCPU limits.

When a VM host object is created (during discovery), the `vmhost.vcpu.max` value is set to the number of physical cores multiplied by a factor of 4. For example, on a Xen VM host that has eight physical cores, the Orchestrate VM host discovery for Xen creates a VM host object with a maximum vCPU of 32. The factor value of 4 represents partitioning a physical core to four vCPUs, which represents 25% capacity.

You can change this default value by creating a policy that sets the `vmhost.vcpu.max` fact value and associates the policy either to the VM host or to a Resource Group of VM hosts. In the preceding Xen VM host example, if you wanted to partition the eight physical cores on a VM host to 50% capacity, you would set the maximum vCPUs to 16, doing so by creating the following policy and then associating it to the VM host or to a Resource Group of VM hosts:

```
<policy>
  <vmhost>
    <fact name="vcpu.max" type="Integer" value="16" />
  </vmhost>
</policy>
```

If you want PlateSpin Orchestrate to allocate vCPUs without checking limits, you can set the maximum vCPUs to -1, which indicates an unlimited number. You would create the following policy to make that configuration setting:

```
<policy>
  <vmhost>
    <fact name="vcpu.max" type="Integer" value="-1" />
  </vmhost>
</policy>
```

6.8.2 Configuring vCPUs on VM Host Clusters

Because the VM host Cluster object represents a set of VM hosts, both `vmhost.vcpu.available` and `vmhost.vcpu.max` facts are sums of the underlying VM host objects. If all of the underlying VM hosts have their `vmhost.vcpu.max` values set to -1 (unlimited), then the corresponding `vmhost.vcpu.max` fact in the VM host Cluster is -1 (unlimited). This also means that no counting or checking occur for vCPU limits. This is used in scenarios where you rely on the underlying hypervisor to account for vCPUs rather than PlateSpin Orchestrate. For example, if you have set up vSphere with DRS and clustering, you do not need PlateSpin Orchestrate to do any checking.

NOTE: The Orchestrate vsphere provisioning adapter sets the `vmhost.vcpu.max` fact value to -1 (unlimited) for DRS-enabled VM host clusters.


6.8.3 Configuring vCPUs on VMs

An existing vCPU fact (`resource.vm.vcpu`) specifies the number of vCPUs for a VM. This fact is set on VM image discovery. You can view it in the Development Client on the VM object Admin view on the Info/Groups page. The Virtual Machine Configuration pane on this page has a *Host CPU Number* field where this fact is set.

In the VM Hosts Monitor view, the tool tip for a VM host displays the following:

```
..Total vCPUs: xx, Available vCPUs: xx
```

Figure 6-1 Tool Tip Text for VM Host vCPU Slots



These values change when a VM is provisioned or shut down. For example, if the `vmhost.vcpu.max` fact were set to -1 (unlimited), then the tool tip would show “-1” for both total and available.

The Virtual Disk Object

7

This section includes the following information:

- ♦ [Section 7.1, “Understanding the Virtual Disk Object,” on page 97](#)
- ♦ [Section 7.2, “Viewing Virtual Disk Configuration in the Development Client,” on page 102](#)

7.1 Understanding the Virtual Disk Object

A virtual disk (vDisk) represents any type of physical disk (such as a file-backed disk image, an ISO image file, a physical hard drive, a block device, a CD/DVD device, or a block device) associated to a VM. The vDisk objects are discovered, along with their associated VM, when a “Discover VM Images” job is run on a repository.

The vDisk is modeled as a Grid object, located as a subordinate to the VM Grid object in the Explorer Tree (of the Development Client). In the Explorer Tree, a vDisk is given the form *vmname_vDisk<n>* where *<n>* represents the numerical order in which this vDisk was discovered, with 1 being appended to the name of the first vDisk discovered or created. For example, *suse11_vdisk1* would be the name of the first disk discovered for a VM with the Grid ID *suse11*. Each additional vDisk is incremented by one, so the second vDisk in this example would be named *suse11_vdisk2*.

This section includes the following information:

- ♦ [Section 7.1.1, “Creating Or Deleting a vDisk in the Development Client,” on page 97](#)
- ♦ [Section 7.1.2, “Sharing Virtual Disks Among VM Hosts,” on page 100](#)
- ♦ [Section 7.1.3, “Moving Virtual Disks,” on page 100](#)

7.1.1 Creating Or Deleting a vDisk in the Development Client

This section includes the following information:

- ♦ [“Creating and Configuring a Virtual Disk” on page 97](#)
- ♦ [“Deleting a Virtual Disk” on page 99](#)

Creating and Configuring a Virtual Disk

You might want to manually create a vDisk in a scenario similar to any of the following:

- ♦ When you want to create a “blank” disk image file for the VM. In this scenario, the disk image does not actually reside on the file system, but a disk image of the specified size (measured in MB) should be created at the location specified for use by the VM.

NOTE: This is essentially a blank file, until used by the VM.

- ♦ When PlateSpin Orchestrate might not have discovered the vDisk objects correctly, such as omitting a disk that should exist. You need to manually correct the incorrect discovery.

- ♦ A VM that already exists needs to have patches applied to it. The patches are delivered through an ISO file, which was not configured to be attached to the VM. This configuration lets the administrator configure the VM with access to the ISO disk image, then apply the patches, and then later delete the vDisk object, returning the VM to its original configuration.

You need to manually add the vDisk, run the Save Config command from the Development Client, and then apply the patches to the running VM. Later, you shut down the VM, delete the vDisk object from PlateSpin Orchestrate, and then run the Save Config command again.

The scenario includes configuring the VM to use the existing ISO file (create vDisk object, running Save Config), and then de-configuring the VM to no longer use the ISO file (delete vDisk object, run saveConfig).

NOTE: At the end of this scenario, only the vDisk object from PSO is deleted, not the ISO file.

To create a virtual disk in the PlateSpin Development Client, you can either right-click the VM where you want to create the vDisk, then select *Create Virtual Disk* (if you do this, you can skip to [Step 4](#) below) or you can use the following procedure from the Development Client menu:

- 1 From the PlateSpin Development Client, select *Actions > Create Virtual Disk* to display the Create a New Virtual Disk dialog box.



- 2 In the *VM* drop-down list, select the name of the VM where you want to add a vDisk, then click *Create*.
- 3 When you have created all of the vDisks you need, click *Close*.
- 4 Select the newly created vDisk object in the Explorer tree to view the Info/Groups page of the Admin view.
- 5 On the Info/Groups page, configure the following settings:
 - ♦ **Type:** Specifies the vDisk [type](#) as the VM host sees it.
 - ♦ **Description:** [Describes](#) the vDisk with any text that you choose
 - ♦ **Healthy:** Designates the [health](#) state of the vDisk.
 - ♦ **Moveable:** Specifies whether the disk image can be copied (relocated) with the VM when the VM is moved (relocated) to another repository. For more information, see [“Moveable”](#) later in this section.
 - ♦ **Repository:** The [repository](#) where this disk location path resides.
 - ♦ **Location:** The path ([location](#)) to the disk image.
 - ♦ If you specify a location to a disk that already exists, the existing disk file is used and the VM configuration is modified accordingly to use this existing disk.


- ♦ If you specify a path to a disk that does not exist, an empty disk image file of the specified size is created.
- ♦ For a vDisk created for a Hyper-V VM, you need to provide the complete path of that vDisk file.

To form the path, you need to know the repository path where the VM currently resides, the vDisk name, which is the name you give it plus the “.vhd” extension. For example, the syntax would be

<value of the “repository.preferredpath” fact>\<your_vhd_filename>.vhd

TIP: Make sure that the .vhd file you designate in this field doesn’t already exist in the path.

- ♦ **Size:** The [size](#) (measured in MB) of the disk image.

6 Click the save icon  to save the fact changes you have made.

7 In the Explorer tree, right-click the VM object where you added the vDisk, then select *Save Config* to apply the changes to the VM’s configuration.

Deleting a Virtual Disk

You might want to manually delete a vDisk in at least two scenarios.

- ♦ When PlateSpin Orchestrate might not have discovered the vDisk objects correctly, such as adding a disk that should not exist. The administrator needs to manually correct the incorrect discovery.
- ♦ A VM that already exists needs to have patches applied to it. The patches are delivered through an ISO file, which was not configured to be attached to the VM. This configuration lets the administrator configure the VM with access to the ISO disk image, then apply the patches, and then later delete the vDisk object, returning the VM to its original configuration.

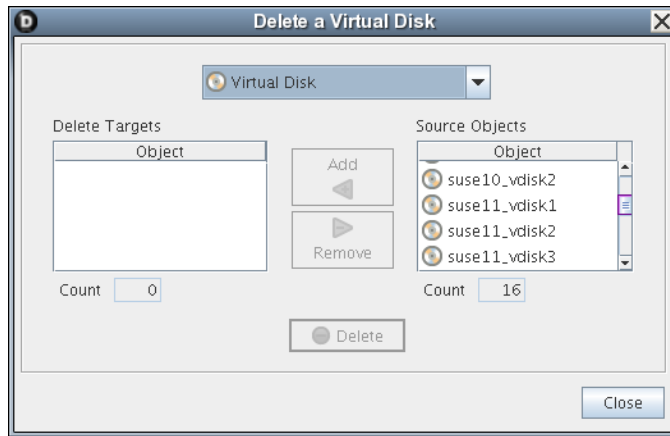
The administrator needs to manually add the vDisk, run the Save Config command from the Development Client, and then apply the patches to the running VM. Later, the administrator shuts down the VM, deletes the vDisk object from PlateSpin Orchestrate, and then runs the Save Config command again.

The scenario includes configuring the VM to use the existing ISO file (create vDisk object, running Save Config), and then de-configuring the VM to no longer use the ISO file (delete vDisk object, run saveConfig).


NOTE: At the end of this scenario, only the vDisk object from PSO is deleted, not the ISO file.

To delete a virtual disk, you can either right-click the vDisk object in the Explorer, then select *Delete* (if you do this, you can skip to [Step 4](#), below), or you can use the following procedure:

- 1** From the PlateSpin Development Client, select the VM object where the vDisk resides, then select *Delete Virtual Disk* to display the Delete a Virtual Disk dialog box.



- 2 In the *Source Objects* list, select the name of the vDisk (hold down the Ctrl key to select multiple), then click *Add* to move these objects to the *Delete Targets* list.
- 3 When you have selected all of the vDisks you want to delete, click *Delete* to display the Delete query dialog box.



- 4 In the dialog box, select *Apply to all* to delete all of the vDisk objects in the *Delete Targets* list, click *OK*, then click *Close*.
- 5 In the Explorer tree, right-click the VM object where you deleted the vDisk, then select *Save Config* to apply the changes to the VM's configuration.

NOTE: The Save Config action rewrites the configuration file for the VM (for example, `config.xen`), but it does not delete any vDisk files on the file system. In this case, manual deletion of the vDisk file would be required.

7.1.2 Sharing Virtual Disks Among VM Hosts


For a VM to be provisionable by other VM hosts, all of a VM's vDisks must be visible in the same way that the VM's default repository (`resource.vm.repository`) is visible to VM hosts. If a VM has multiple vDisks and each vDisk has a different associated repository, these repositories must also be visible from a potential VM host.

7.1.3 Moving Virtual Disks

When you move a VM to a new repository, all of its moveable vDisk images (see [“Moveable:” on page 103](#)) are moved with it to be co-located in the same repository. PlateSpin Orchestrate uses the aggregated size of each moveable vDisk to determine if the designated repository has enough space for all of the disk images. vDisks that are marked as not moveable stay in place and are not used in the calculation for the VM disk size.

The following illustration further explains this concept:

Figure 7-1 Example of Moving Virtual Disks with the VM



- VM host 1, VM host 2, and VM host 3 all have their own local storage repositories.
- VM host 1 has a vDisk located on it. It is designated as a moveable vDisk.
- VM host 1 and VM host 2 are also connected to a shared NAS storage repository.
- The local repository connected to VM host 1 has a vDisk located on it. It is designated as a moveable vDisk.
- The shared NAS repository has a vDisk located on it. It is designated as a non-moveable vDisk.

NOTE: Shared repositories are not created on discovery. They must be manually created and the sharing (visibility) configured.

- VM host 1 has a VM located on it.
- VM host 3 cannot communicate with the NAS repository; its `vmhost.repositories` fact does not include the NAS repository in the array, so that repository is not visible to VM host 3.

If you want to move the VM from VM host 1 to another VM host, PlateSpin Orchestrate manifests the following behavior:

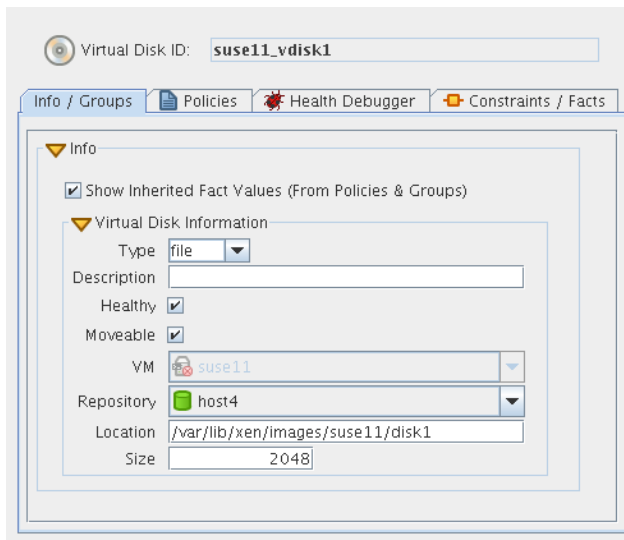
- The vDisk sizes used by the VM (on local storage and shared storage) are aggregated and compared to free space available on the repositories
- The only vDisk that would be allowed to copy is the moveable disk. This disk would be copied to either the shared NAS repository or the local storage on VM host 2.
- VM host 3 would not be considered because it does not have access to the unmoveable disk on the NAS repository.

7.2 Viewing Virtual Disk Configuration in the Development Client



You can visually expose a vDisk Grid object  in the Development Client in two ways:

- ♦ In the Explorer Tree, select a VM Resource object, then select the *Info/Groups* tab in the Admin View to open the VM Info Groups page, then scroll to the *Virtual Machine Configuration* panel on that page. You can right-click the vDisk icon in that panel to display the four tabs in the Virtual Disk Admin view.
- ♦ In the Explorer Tree, click the expand/collapse icon of a VM Grid object, identify the vDisk icon, then select the icon to display the four tabs in the Virtual Disk Admin view.

Figure 7-2 The Virtual Disk Info/Groups Page



The page that opens under the *Info* tab includes fields where you can configure the general information and attributes (facts) of the vDisk.

NOTE: Whenever you make changes to vDisk object facts, the write icon is superimposed on the object's icon , signifying that the object has been altered. If you want to save the changes you have made, you need to click the *Save* icon  on the Development Client toolbar.

This section includes the following additional information:

- ♦ [Section 7.2.1, “The Virtual Disk Information Panel,” on page 102](#)
- ♦ [Section 7.2.2, “The Virtual Disk Policies Tab,” on page 104](#)
- ♦ [Section 7.2.3, “The Virtual Disk Health Debugger Tab,” on page 104](#)
- ♦ [Section 7.2.4, “The Virtual Disk Constraints/Facts Tab,” on page 105](#)
- ♦ [Section 7.2.5, “Virtual Disk Object Naming and Renaming,” on page 105](#)

7.2.1 The Virtual Disk Information Panel

The Virtual Disk Information panel on the Info page includes the following fields:

NOTE: Tooltip text is available when you mouse over any of these fields.

Type: This drop down list lets you select one of the vDisk types as the VM host sees it:

- ♦ **file:** Specifies this vDisk as a file-backed disk.
- ♦ **block:** Specifies this vDisk as a block device.

In the Fact Editor, this fact is listed as `vdisk.type`:

```
<fact name="vdisk.type" value="file" type="String" />
```

Description: Describes the vDisk with any text that you choose.

In the Fact Editor, this fact is listed as `vdisk.description`:

```
<fact name="vdisk.description" value="" type="String" />
```

For a vDisk discovered and managed by the Xen provisioning adapter, this field is usually blank. Because this is a freeform field, you can enter any text you want here. For a vDisk discovered and managed by the vSphere provisioning adapter, this field is populated with a display name label obtained by vSphere and mapped to this vDisk by PlateSpin Orchestrate.

Healthy: For a vDisk managed by Xen, this check box is selected by default, which designates the vDisk as being in good health.

NOTE: We recommend that you do not change the *Healthy* value from its default.

In the Fact Editor, this fact is listed as `vdisk.health`:

```
<fact name="vdisk.health" value="true" type="Boolean" />
```

Moveable: When this check box is selected (its value is true), the vDisk is moveable, which means that the disk image can be copied to a different repository when the VM moves.

In the Fact Editor, this fact is listed as `vdisk.moveable`:

```
<fact name="vdisk.moveable" value="true" type="Boolean" />
```

If *Moveable* is unchecked, the disk image must stay at its current location because it cannot be copied or moved. By default upon discovery, if PlateSpin Orchestrate sees that this vDisk is an ISO image, the fact is set to false because it is assumed that the administrator doesn't want to copy ISO images from one location to another.

Whenever you want to prevent a vDisk from being moved, you can deselect this check box.

VM: (Read Only) Specifies the name of the VM that uses this vDisk.

In the Fact Editor, this fact is listed as `vdisk.vm`:

```
<fact name="vdisk.vm" value="mysql" type="String" />
```

This is a fact junction referencing the associated VM. Conversely, the `resource.vm.vdisks` fact visible from the VM Grid object is a fact junction showing the associated vDisks associated with the VM.

Repository: The storage location containing the vDisk image on the VM host. In the case of block type disks, the repository should be typically set to `--none--`, which means that it will pass all of the repository constraints, so this setting would not take effect when choosing a host.

NOTE: Changing this fact after discovery only corrects a possible incorrectly discovered fact. Changing the storage location does not move the vDisk.

In the Fact Editor, this fact is listed as `vdisk.repository`:

```
<fact name="vdisk.repository" value="zos" type="String" />
```

Location: For file-backed disks, this fact represents the file system path to the vDisk image in the specified repository.

For example, a vDisk located on an NFS repository datastore would show the URI to the NFS share with the path to the disk appended to it.

In the case of block type disks, this fact contains the URI to the block device, for example `/dev/hdc`, which could represent a CD/DVD tray on a VM host.

In the Fact Editor, this fact is listed as `vdisk.location`:

```
<fact name="vdisk.location" value="/var/lib/xen/images/mysql/disk1" type="String" />
```

Size: The size (measured in MB) of this vDisk image.

In the Fact Editor, this fact is listed as `vdisk.size`:

```
<fact name="vdisk.size" value="2048" type="Integer" />
```

The disk size value for each moveable vDisk on a VM is aggregated by PlateSpin Orchestrate into the `resource.vm.vdisksize` fact, which is used to determine if the VM can relocate from one repository to another, given that the new repository has enough free space to store the VM.

7.2.2 The Virtual Disk Policies Tab

The *Policies* tab opens a page that contains a policy viewer for each of the policies associated with a Grid object.

NOTE: You can edit a policy by right-clicking a policy icon, selecting *Edit Policy* and clicking the Save icon.

7.2.3 The Virtual Disk Health Debugger Tab

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Chapter A.3, “The Health Debugger,” on page 187](#).

7.2.4 The Virtual Disk Constraints/Facts Tab

To support constraining a VM's provision actions based on more than one disk's repository (that is, more than just `resource.vm.repository`), the vDisk can be referenced in constraints. The vDisk constraints are used to assign VM hosts during actions such as provisioning, building, or migrating. You can write constraints against attributes of disks (such as the repository where the vDisk resides) and against the available VM host repositories.

The Constraints/Facts tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. In essence, by building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any and all system resources by managing an object's facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode `r/o` have read-only values, which can be viewed (using the edit pencil icon) but changes cannot be made.

IMPORTANT: Several custom facts for the vDisk object can be added at discovery time, and vary according to the VM technology that manages the respective vDisk. These facts are documented in the [PlateSpin Orchestrate 2.5 Developer Guide and Reference](#).

For more information about using the Fact Editor on this page to rename the Virtual Disk object, see [Section 7.2.5, “Virtual Disk Object Naming and Renaming,” on page 105](#).

7.2.5 Virtual Disk Object Naming and Renaming

Some resource names are generated by the Orchestrate system and can therefore receive generic, arbitrary names such as `mysql-vdisk1`, `mysql-vdisk2`, and so on. A Virtual Disk (vDisk) you name at creation time might also change later in its purpose or facilities.

As the quantity of these vDisk objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object's “display name” is visible in the Development Client interface, the VM Client interface, the Server Portal, and in optional `zos` and `zosadmin` commands.

NOTE: Resource object groups (that is, the folders that contain these vDisk objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A vDisk object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a vDisk object in the Orchestrate Development Client using one of three methods:

- ♦ Right-click the vDisk object in the Explorer tree, then select **Rename** to allow editing of the display name.
- ♦ Triple-click the vDisk object in the Explorer tree to allow editing of the display name.
- ♦ In the *Constraints/Facts* page, where you can select the vDisk object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you begin to rename using one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the vDisk object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the `zos` or `zosadmin` command line, see the [PlateSpin Orchestrate 2.5 Command Line Reference](#).

The Virtual NIC Object

8

This section includes the following information:

- ♦ [Section 8.1, “Understanding the Virtual NIC Object,” on page 107](#)
- ♦ [Section 8.2, “Viewing the Virtual NIC Configuration in the Development Client,” on page 110](#)

8.1 Understanding the Virtual NIC Object

A virtual network interface card (vNIC) represents the configuration of a VM connected to a network. A VM can be configured to have multiple vNICs. When a VM is provisioned, each of its associated vNICs can be attached to a [Virtual Bridge](#) in order to gain connectivity to a specified network. The vNIC objects are discovered, along with their associated VM, when a “Discover VM Images” job has been run on a repository.

The vNIC is modeled as a Grid object, located as a subordinate to the VM Grid object in the Explorer Tree (of the Development Client). A vNIC is given the form of *vmname_vnic<n>* where *<n>* is appended to indicate the order of discovery or creation of the vNIC. For example, *redhat_vnic1* would be the name of the first NIC discovered for a VM with the Grid ID *redhat*. Each additional vNIC would be incremented by one, so the second vNIC in this example would be named *redhat_vnic2*.

This section includes the following information:

- ♦ [Section 8.1.1, “The Purpose of the Virtual NIC,” on page 107](#)
- ♦ [Section 8.1.2, “Creating Or Deleting a vNIC in the Development Client,” on page 107](#)

8.1.1 The Purpose of the Virtual NIC

A vNIC represents the network interface configuration for a virtual machine. A vNIC is linked to a network by connecting to a virtual network bridge (vBridge). A group of vBridge objects is represented as a Network group in the Explorer Tree. By convention, during VM host discovery, any vBridges that are configured with the same name are assumed to be part of the same network.

For more information, see [Chapter 9, “The Network Group and its Virtual Bridge Objects,” on page 117](#).

8.1.2 Creating Or Deleting a vNIC in the Development Client

Although a vNIC is generally discovered on a VM, you can also manually create or delete a vNIC. This section includes the following information:

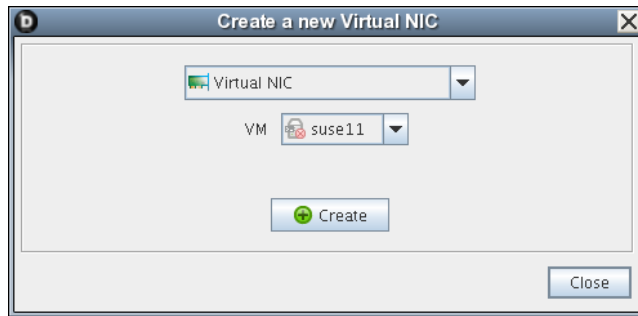
- ♦ [“Creating a Virtual NIC” on page 108](#)
- ♦ [“Creating a Virtual NIC for a Hyper-V VM” on page 108](#)
- ♦ [“Deleting a Virtual NIC” on page 109](#)

Creating a Virtual NIC

You might want to manually create a vNIC anytime you want to give a VM access to a network configured on the VM host.


To create a vNIC, you can select the VM Grid object in the Explorer tree, then right-click and select *Create Virtual NIC*. You can also use the following alternate method:

- 1 From the PlateSpin Development Client, select *Actions > Create Virtual NIC* to display the Create a New Virtual NIC dialog box.



- 2 In the VM drop-down list, select the name of the VM to which you want to add a vNIC, then click *Create*.
- 3 When you have created all of the vNICs you need, click *Close*.
- 4 Select the newly created vNIC object in the Explorer Tree to view the Info/Groups page of the admin view.
- 5 On the Info/Groups page, configure the following settings:
 - ♦ **MAC Address:** The [MAC Address](#) assigned to this vNIC. If left empty, or set to asterisk (*), a MAC address is autogenerated for this vNIC.
 - ♦ **Network:** The network (vBridge group) that should be used when provisioning the VM. When a VM is provisioned, any vBridge contained within the specified network group can be used for attaching the vNIC to the network.


Although the vNIC can be formally created at this point, you can also configure [Autoprep](#) or [Sysprep facts](#) used to prepare a personalized version of the VM that can be provisioned later.

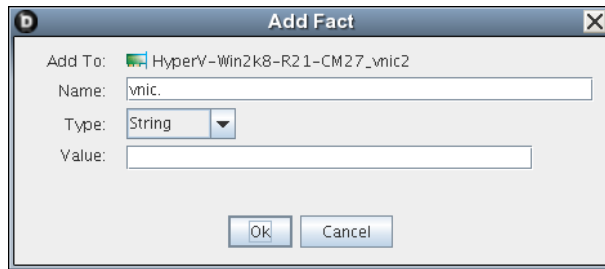
- 6 Click the save icon  to save the fact changes you have made.
- 7 In the Explorer tree, right-click the VM object to which you added the vNIC, then select *Save Config* to apply the changes to the VM's configuration.

IMPORTANT: If you do not run the *Save Config* action after configuring the vNIC, any vNICs that were added to the grid or vNIC settings that were modified could be lost.

Creating a Virtual NIC for a Hyper-V VM

Use the following steps to create a vNIC for a Hyper-V VM:

- 1 Right-click the VM for which you want to create a vNIC, select *Create Virtual NIC*.
- 2 In the Explorer tree, select the vNIC object you created in the previous step, then from the Admin view, select the *Constraints/Facts* tab to open the Constraints/Facts page.
- 3 In the Constraints/Facts page, click the add a fact  button to open the Add Fact dialog box.



- 4 In the dialog box, add a custom fact for each of the items in the table below.

Fact name	Type	Value and Description
vnic.type	String	SyntheticEthernetPort or EmulatedEthernetPort SyntheticEthernetPort is the default vNIC type. Use the EmulatedEthernetPort type to perform a network-based installation of the guest OS or when integration services are not installed in the guest OS.
vnic.mac	String	Provide a valid MAC address.
vnic.static_mac_address	Boolean	true or false. When set to true, the provided MAC address (vnic.mac) is used to set for this adapter, otherwise it is dynamically set by the Hyper-V system.
vnic.network	String	Copy the value of group.id fact from the networks to which you want to attach the vNIC.

- 5 In the Explorer tree, right-click the VM object to which you added the vNIC, then select *Save Config* to apply the changes to the VM's configuration.


IMPORTANT: If you do not run the *Save Config* action after configuring the vNIC, any vNICs that were added to the grid or vNIC settings that were modified could be lost.

Deleting a Virtual NIC


Although uncommon, you might want to manually delete a vNIC when you no longer want the VM to have access to a specified network on the VM host. For example, if the initial need for connecting the VM to a network no longer exists or the network is going to become private and the VM should not have access you would delete the virtual NIC that allows connectivity.

To delete a virtual NIC, you can select the vNIC object in the Explorer tree, then right-click and select *Delete*, or you can use the following procedure:

- 1 From the PlateSpin Development Client, select *Actions> Delete Virtual NIC* to display the Create a New Virtual NIC dialog box.



- 2 In the *Source Objects* list, select the name of the vNIC (hold down the Ctrl key to select multiple), then click *Add* to move these objects to the *Delete Targets* list.
- 3 When you have selected all of the vNICs you want to delete, click *Delete* to display the Delete query dialog box.



- 4 In the dialog box, select *Apply to all* to delete all of the vNIC objects in the *Delete Targets* list, click *OK*, then click *Close*.
- 5 In the Explorer tree, right-click the VM object where you deleted the vNIC, then select *Save Config* to apply the changes to the VM's configuration.

NOTE: You must run the *Save Config* action to confirm the deletion of the vNIC. If you do not run this action, the vNIC is not deleted from the VM's configuration.

8.2 Viewing the Virtual NIC Configuration in the Development Client

You can visually expose a vNIC Grid object  in the Development Client in two ways:

- ♦ In the Explorer Tree, select a VM Resource object, then select the *Info/Groups* tab in the Admin View to open the VM Info Groups page, then scroll to the *Virtual Machine Configuration* panel on that page. You can right-click the vNIC icon in that panel to display the four tabs in the Virtual Disk Admin view.
- ♦ In the Explorer Tree, click the expand/collapse icon of a VM Grid object, identify the vNIC icon, then select the icon to display the four tabs in the Virtual NIC Admin view.

Figure 8-1 The Virtual NIC Info/Groups Page

Virtual NIC ID: `suse11_vnic2`

Info / Groups Policies Health Debugger Constraints / Facts

Info

☒ Show Inherited Fact Values (From Policies & Groups)

Virtual NIC Information

Description

Healthy ☒

MAC Address

Vbridge --none--

VM suse11

Network --none--

Autoprep Network Adapter

MAC Address < undefined > Define



Use DHCP < undefined > Define

IP Address < undefined > Define

Subnet Mask < undefined > Define

DNS from DHCP < undefined > Define

The page that opens under the *Info/Groups* tab includes fields where you can configure the general information and attributes (facts) of the vNIC.

NOTE: Whenever you make changes to vNIC object facts, the write icon is superimposed on the object's icon , signifying that the object has been altered. If you want to save the changes you have made, you need to click the *Save* icon  on the Development Client toolbar.

This section includes the following additional information:

- ♦ [Section 8.2.1, “The Virtual NIC Info Panel,” on page 111](#)
- ♦ [Section 8.2.2, “The Virtual NIC Policies Tab,” on page 115](#)
- ♦ [Section 8.2.3, “The Virtual NIC Health Debugger Tab,” on page 115](#)
- ♦ [Section 8.2.4, “The Virtual NIC Constraints/Facts Tab,” on page 115](#)
- ♦ [Section 8.2.5, “Virtual NIC Object Naming and Renaming,” on page 115](#)

8.2.1 The Virtual NIC Info Panel

The Virtual NIC Info panel on the Info/Groups page includes the following sections:

- ♦ [“Show Inherited Fact Values Check Box” on page 111](#)
- ♦ [“Virtual NIC Information” on page 112](#)
- ♦ [“Autoprep/Sysprep Network Adapter” on page 113](#)

Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached and/or inherited policies. Such fact values are read only (non-editable).

Virtual NIC Information

The Virtual NIC Information panel on the Info page includes the following fields:

NOTE: Tool tip text is available when you mouse over any of these fields.

Description: A free-form field you can use to add any description about this vNIC.

In the Fact Editor, this fact is listed as `vnic.description`:

```
<fact name="vnic.description" value="" type="String" />
```

Healthy: In most cases, this check box is selected by default, which designates the vNIC as being in good health.

NOTE: We recommend that you do not change the *Healthy* value from its default.

In the Fact Editor, this fact is listed as `vnic.health`:

```
<fact name="vnic.health" value="true" type="Boolean" />
```

MAC Address: Specifies the MAC address assigned to this vNIC. An empty string implies an auto-generated MAC address, as does an asterisk (*).

When the VM appears on the network, this will be its MAC identifier. A MAC address must be unique on the network to avoid routing conflicts.

There are some situations when you might want to define a static MAC address. For example, if a VM uses DHCP, you might want the DHCP service on the network to give the VM a static address. When the VM boots up and attempts to get an IP address, it contacts the DHCP server, which sees its statically-defined MAC address and then provides the same IP address (not a new one) each time the VM boots up. In this way, the VM is consistently configured with the same IP address. You might also want to define a static MAC address for audit trails or other security reasons.

In the Fact Editor, this fact is listed as `vnic.mac`:

```
<fact name="vnic.mac" value="" type="String" />
```

vBridge: Specifies the name of the actual host bridge used by this vNIC. When the VM is not running, this field is blank. When a VM is provisioned, a vBridge is chosen for this vNIC based on the available VM hosts and the specified network group. When the VM is running, the associated vBridge is identified in this field. For more information, see [Chapter 9, “The Network Group and its Virtual Bridge Objects,” on page 117](#).

In the Fact Editor, this read-only fact is listed as `vnic.vbridge`:

```
<fact name="vnic.vbridge" value="" type="String" />
```

VM: Specifies the name of the VM resource that uses this vNIC .

In the Fact Editor, this read-only fact is listed as `vnic.vm`:

```
<fact name="vnic.vm" value="" type="String" />
```

This is a fact junction referencing the associated VM. Conversely, the `resource.vm.vnics` fact visible from the VM Grid object is a fact junction showing the associated vNICs associated with the VM. A vNIC cannot be shared between two VMs. Each VM has its own vNIC objects.

Network: The network (vBridge group) that should be used by this vNIC when the VM is provisioned. This fact will be used in combination with the VM host placement constraints to choose a suitable vBridge at provision time.

In the Fact Editor, this fact is listed as `vmnic.network`:

```
<fact name="vmnic.network" value="eth1" type="String" />
```

Autoprep/Sysprep Network Adapter

VMs can be prepared for provisioning by configuring the facts in this panel. Click *Define* on each field if the value has not been previously configured.

NOTE: When you change any of the settings in this panel, you need to right click the VM and select *Personalize* for the changes to take effect.

- ♦ **MAC Address:** The MAC address of the interface. Specify an asterisk (*) to generate a new MAC address. If the value is not set, the existing `vmnic.mac` is used.

IMPORTANT: An unset *MAC Address* fact generates a new MAC address. This is contrary to the current tool tip text.

In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.MACAddress`:

```
<fact name="vmnic.provisioner.autoprep.MACAddress" value="" type="String" />
```

- ♦ **Use DHCP:** When this check box is selected (it has a value of true), the VM will be configured to retrieve its network settings from a DHCP server. If the check box is not selected (it has value of false), you should make sure that the IP address, subnet mask, and gateway address facts are defined. In the Fact Editor, this fact is listed as

`vmnic.provisioner.autoprep.UseDHCP`:

```
<fact name="vmnic.provisioner.autoprep.UseDHCP" value="false" type="Boolean" />
```

- ♦ **IP Address:** The IP address for the adapter.

In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.IPAddress`:

```
<fact name="vmnic.provisioner.autoprep.IPAddress" value="" type="String" />
```

- ♦ **Subnet Mask:** The subnet mask for this adapter.


In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.subnetMask`:

```
<fact name="vmnic.provisioner.autoprep.subnetMask" value="" type="String" />
```

- ♦ **Gateway IP Addresses:** (Windows Sysprep only) A list of the gateway IP addresses available to the interface.

In the Fact Editor, this fact is listed as an array:

```
<fact name="vmnic.provisioner.autoprep.Gateways">
  <array type="String">
    </array>
</fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove the IP address or change its order in the array of element choices.

- ♦ **DNS from DHCP:** (Optional. SUSE VM only) When this check box is selected (it has a value of true), the SUSE VM is configured to retrieve its DNS server settings from DHCP.

In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.DNSFromDHCP`:

```
<fact name="vmnic.provisioner.autoprep.DNSFromDHCP" value="false"
type="Boolean" />
```

- ♦ **DNS Server IP Addresses:** (Windows VM only) The adapter's list of DNS servers used for name lookup.

In the Fact Editor, this fact is listed as an array:

```
<fact name="vmnic.provisioner.autoprep.DNSServers">
  <array type="String">
    </array>
  </fact>
```

- ♦ **DNS Domain:** (Windows VM only) The adapter's DNS domain name.

In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.DNSDomain`:

```
<fact name="vmnic.provisioner.autoprep.DNSDomain" value="" type="String" />
```

- ♦ **Primary WINS Server:** (Windows VM only) The name of the adapter's primary WINS server.

In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.primaryWINS`:

```
<fact name="vmnic.provisioner.autoprep.primaryWINS" value="" type="String"
/>
```

- ♦ **Secondary WINS Server:** (Windows VM only) The name of the adapter's secondary WINS server.


In the Fact Editor, this fact is listed as `vmnic.provisioner.autoprep.secondaryWINS`:

```
<fact name="vmnic.provisioner.autoprep.secondaryWINS" value=""
type="String" />
```

- ♦ **DNS Suffixes:** (Windows VM only) A list of the suffixes associated with this adapter that are appended to the name for lookup.

In the Fact Editor, this fact is listed as an array:

```
<fact name="vmnic.provisioner.autoprep.DNSSuffixes">
  <array type="String">
    </array>
  </fact>
```

You can edit this array by clicking the  button to open an array editor. In this dialog box you can add or remove the DNS suffix or change its order in the array of element choices.

- ♦ **NetBIOS:** (Windows VM only) The NetBIOS options for this VM. Options include:

- ♦ *EnableNetBIOSviaDhcp*
- ♦ *EnableNetBIOS*
- ♦ *DisableNetBIOS*

In the Fact Editor this fact is listed as `vmnic.provisioner.autoprep.netBIOS`:

```
<fact name="vmnic.provisioner.autoprep.netBIOS" value="" type="String" />
```

8.2.2 The Virtual NIC Policies Tab

The *Policies* tab opens a page that contains a policy viewer for each of the policies associated with a Grid object.

NOTE: You can edit a policy by right-clicking a policy icon, selecting *Edit Policy* and clicking the Save icon.

8.2.3 The Virtual NIC Health Debugger Tab

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Chapter A.3, “The Health Debugger,” on page 187](#).

8.2.4 The Virtual NIC Constraints/Facts Tab

The Constraints/Facts tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. In essence, by building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any and all system resources by managing an object’s facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode *r/o* have read-only values, which can be viewed (that is, using the edit “pencil” icon) but changes cannot be made.

For information about using the Fact Editor on this page to rename the Virtual NIC object, see [Section 8.2.5, “Virtual NIC Object Naming and Renaming,” on page 115](#).

8.2.5 Virtual NIC Object Naming and Renaming

Some resource names are generated by the Orchestrate system and can therefore receive generic, arbitrary names such as `mysql-vnic1`, `mysql-vnic2`, and so on. A Virtual NIC (vNIC) you name at creation time might also change later in its purpose or facilities.

As the quantity of these vNIC objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object’s “display name” is visible in the Development Client interface, the VM Client interface, the Server Portal, and in optional `zos` and `zosadmin` commands.

NOTE: Resource object groups (that is, the folders that contain these vNIC objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A vNIC object’s name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a vNIC object in the Orchestrate Development Client using one of three methods:

- ♦ Right-click the vNIC object in the Explorer tree, then select *Rename* to allow editing of the display name.

- ♦ Triple-click the vNIC object in the Explorer tree to allow editing of the display name.
- ♦ In the *Constraints/Facts* page, where you can select the vNIC object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you begin to rename using one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the vNIC object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see the [PlateSpin Orchestrate 2.5 Command Line Reference](#).

The Network Group and its Virtual Bridge Objects

9

This section includes the following information:

- [Section 9.1, “Understanding the Network Group and Virtual Bridge Objects,” on page 117](#)
- [Section 9.2, “Viewing the Virtual Bridge Configuration in the Development Client,” on page 121](#)

9.1 Understanding the Network Group and Virtual Bridge Objects

- [Section 9.1.1, “The Virtual Bridge Object,” on page 117](#)
- [Section 9.1.2, “The Purpose of the Virtual Bridge,” on page 118](#)
- [Section 9.1.3, “Creating or Deleting a vBridge in the Development Client,” on page 118](#)
- [Section 9.1.4, “Virtual Bridge Object Naming and Renaming,” on page 120](#)

9.1.1 The Virtual Bridge Object

In PlateSpin Orchestrate, a group of Virtual Bridge (vBridge) objects is called a Network. It represents the networks (actually, LANs and VLANs) that are available to that VM host, and can be shared across multiple VM hosts. Network groups are created automatically during VM and VM host discovery as their virtual networking settings are determined. By default, PlateSpin Orchestrate automatically groups vBridges with the same name into a corresponding Network with the same name, assuming that vBridges with similar names usually refer to the same Network. After discovery, you can freely reassign the vBridges to other Networks or multiple Networks, depending on your organization’s physical network topology.

The VLAN Fact


A Virtual LAN (VLAN) provides a “logical” LAN topology for a group of machines that does not have to depend on the switch hardware to which the machines (virtual or physical) are directly connected. Modern “smart” network switches can keep track of a “VLAN ID” on network traffic passing through them. Using this ID, the switches transparently route such traffic to the hosts configured to use the VLAN specified by the ID.

Using VLANs can reduce the costs of an organization’s physical network infrastructure. For example, if your organization requires multiple subnets or multi-homed hosts, you won’t have to buy new equipment for each new subnet or install multiple physical NICs on the physical machines; this can be done with VLANs using the same common physical network layer. VLANs allow greater flexibility in managing the network topology.

Where possible, PlateSpin Orchestrate discovers that a VLAN ID is in use on a network. The Virtual LAN Identifier fact is found on a Network object :

```
<fact name="group.vlanid" value="" type="String" />
```

The fact is populated with a positive integers value upon discovery of an already-existing VM host configuration. PlateSpin Orchestrate can track a VLAN ID for each network object to allow correct management of individual VLANs as full-fledged networks.

PlateSpin Orchestrate applies a graphic overlay to the Network icon  to signify that the Network was discovered as a VLAN. In cases where VLAN discovery is not accomplished through automation, you can also set the VLAN fact value by hand.

9.1.2 The Purpose of the Virtual Bridge

The vBridge is discovered, along with its associated VM, when a discovery job runs on a VM host. A virtual bridge (vBridge) acts as a “virtual” Ethernet segment contained entirely within the software of a VM host. The virtual NIC (vNIC) devices on that host’s VMs can each be assigned to one of the VM host’s vBridges as if they were being physically connected to a LAN.

Virtual bridges can also be associated with one or more physical NICs to combine the virtual LANs on these hosts into one common virtual LAN. This combined LAN is referred to as a “Network” in PlateSpin Orchestrate. Association of virtual bridges is also frequently done to include a “virtual” LAN on a VM host in the organization’s overall physical LAN topology so the VMs can access other systems in the organization as if they were directly connected to the switches.

The following points might help you understand the relationship of these objects:

- ♦ Associating a vNIC to a vBridge is like plugging a physical NIC into an Ethernet switch.
- ♦ Associating a vBridge with a physical NIC is like connecting two physical switches together using their uplink ports.
- ♦ A vBridge can become part of a vLAN by associating it with a physical NIC device that itself is configured as a VLAN.
- ♦ A VM host can have multiple vBridges, each of which can be connected to separate physical networks (for example, through 802.1Q VLAN tagging).
- ♦ When a VM is provisioned, a Virtual NIC must be connected to a Virtual Bridge in order for the Virtual NIC to be usable.

In the Explorer Tree, a vBridge is given the form *vmhostname_ethn* where *n* represents the numerical order in which this vBridge was discovered on the VM host, with 0 being appended to the name of the first vBridge discovered or created. For example, *host1_eth0* might be the name of the first bridge. Each additional vBridge is incremented by one, so the second vBridge in this example would be named *host1_eth1*.

9.1.3 Creating or Deleting a vBridge in the Development Client

This section includes the following information:

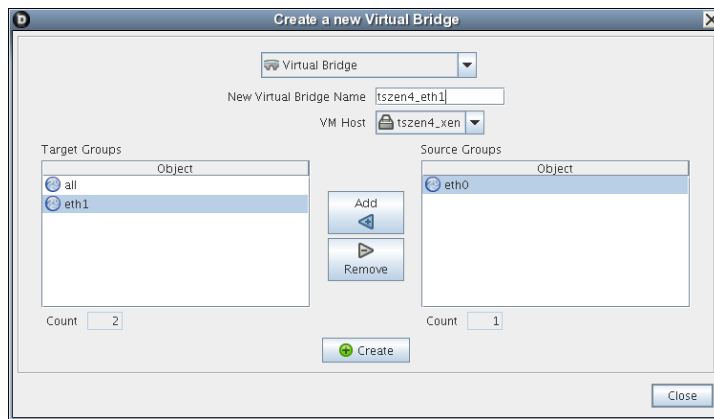
- ♦ [“Creating a Virtual Bridge” on page 119](#)
- ♦ [“Deleting a Virtual Bridge” on page 119](#)

Creating a Virtual Bridge

You might want to manually create a vBridge if, for some reason, the discovery process did not find one of the vBridges in a network. Creating a new vBridge in the Development Client does not “add” a physical bridge to the network, it only helps to model a physical bridge that was not previously discovered.


To create a vBridge, you can select the Network object in the Explorer tree, then right-click and select *New Virtual Bridge*. You can also use the following alternate method:

- 1 From the PlateSpin Development Client, select *Actions > Create Virtual Bridge* to display the Create a New Virtual Bridge dialog box.



- 2 In the *Source Groups* list, select the Network object where you want to add a vBridge, then click *Add* to move it to the *Target Groups* list.
- 3 In the *New Virtual Bridge Name* field, enter the name you want to use to identify this vBridge. The data you enter here is completely free form. When PlateSpin Orchestrate discovers and names vBridges, it associates them by name as the default practice. This is done for convenience in locating the objects. A vBridge can be named anything, provided that it is a legal name for a vBridge on the VM host's operating system.

It is not necessary to add the `<vmhostname>_` prefix on the name. This prefix is automatically prepended when the new object is created. For example, if you select bridge name `sales` on VM host `vmh1`, the actual object is created as `vmh1_sales`.


- 4 Click *Create*, then click *Close*.
- 5 In the Explorer tree, expand the Network group where you created the new vBridge object.
- 6 Select the new vBridge object to open its Info/Groups admin view.
- 7 Configure the settings described in [Section 9.2.1, “The Virtual Bridge Info/Groups Tab,”](#) on [page 122](#).
- 8 Click the save icon  to save the fact changes you have made.

Deleting a Virtual Bridge


Although uncommon, you might want to manually delete a vBridge when you no longer want the VM to have access to a specified network on the VM host. For example, if the initial need for connecting the VM to a network no longer exists or the network is going to become private and the VM should not have access you would delete the vBridge that allows connectivity.

To delete a Virtual Bridge, you can select the vBridge object in the Explorer tree, then right-click and select *Delete*, or you can use the following procedure:

- 1 From the PlateSpin Development Client, select *Actions > Delete Virtual Bridge* to display the Create a New Virtual Bridge dialog box.



- 2 In the *Source Objects* list, select the name of the vBridge (hold down the Ctrl key to select multiple), then click *Add* to move these objects to the *Delete Targets* list.
- 3 When you have selected all of the vBridges you want to delete, click *Delete* to display the Delete query dialog box.



- 4 In the dialog box, select *Apply to all* to delete all of the vBridge objects in the *Delete Targets* list, click *OK*, then click *Close*.

9.1.4 Virtual Bridge Object Naming and Renaming

Some resource names are generated by the Orchestrate system and can therefore receive generic, arbitrary names such as *host1-eth1*, *host2-eth1*, and so on. A Virtual Bridge (vBridge) you name at creation time might also change later in its purpose or facilities.

As the quantity of these vBridge objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object's "display name" is visible in the Development Client interface, the VM Client interface, the Server Portal, and in optional *zos* and *zosadmin* commands.

NOTE: A Network object (that is, the group that contains these vBridge objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A vBridge object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a vBridge object in the Orchestrate Development Client using one of three methods:

- ♦ Right-click the vBridge object in the Explorer tree, then select *Rename* to allow editing of the display name.

- Triple-click the vBridge object in the Explorer tree to allow editing of the display name.
- In the *Constraints/Facts* page, where you can select the vBridge object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you begin to rename using one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the vBridge object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the `zos` or `zosadmin` command line, see the [PlateSpin Orchestrate 2.5 Command Line Reference](#).

9.2 Viewing the Virtual Bridge Configuration in the Development Client


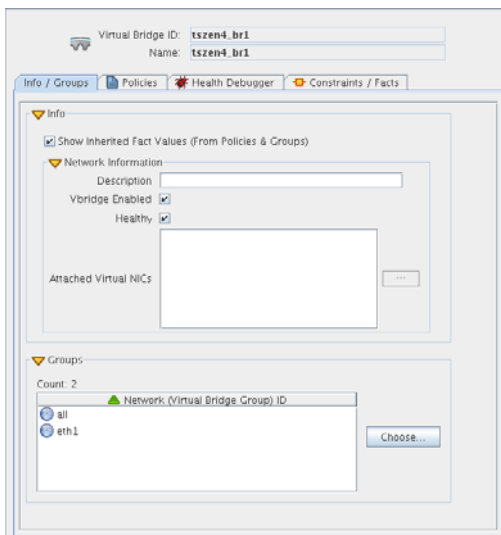


You can visually expose a vBridge Grid object  in the Development Client Explorer tree by selecting a Network object to expand it, then selecting the vBridge you want to view in the Info/Groups page.

Figure 9-1 The Virtual Bridge Info/Groups Page



The page that opens under the *Info/Groups* tab includes fields where you can configure the general information and attributes (facts) of the vBridge.

NOTE: Whenever you make changes to vBridge object facts, the write icon is superimposed on the object's icon , signifying that the object has been altered. If you want to save the changes you have made, you need to click the *Save* icon  on the Development Client toolbar.

9.2.1 The Virtual Bridge Info/Groups Tab

The page that opens under the *Info/Groups* tab includes two collapsible sections where you can configure the general information and attributes of the vBridge.

- ♦ [“The Virtual Bridge Info Panel” on page 122](#)
- ♦ [“The Virtual Bridge Groups Panel” on page 123](#)

The Virtual Bridge Info Panel

The *Info* panel on the *Info/Groups* page includes the following information:

- ♦ [“Show Inherited Fact Values Check Box” on page 122](#)
- ♦ [“Network Information” on page 122](#)

Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached and/or inherited policies. Such fact values are read only (non-editable).

Network Information

The *Network Information* panel on the *Info/Groups* page for the vBridge Grid object includes the following fields:

NOTE: Tool tip text is available when you mouse over any of these fields.

Description: Enter a description of the vBridge Grid object.

In the Fact Editor, this fact is listed as `vmhost.resource`:

```
<fact name="vbridge.description" value="" type="String" />
```

Vbridge Enabled: This check box is selected by default. When it is selected (it has a value of true), the vBridge is enabled and Virtual NICs can be attached to it.

In the Fact Editor, this fact is listed as `vbridge.enabled`:

```
<fact name="vbridge.enabled" value="true" type="Boolean" />
```

Healthy: This check box is selected by default. When it is selected (it has a value of true), the Virtual Bridge is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy (this follows the normal rules of policy inheritance). For more information, see [Appendix A, “Grid Object Health Monitoring,” on page 185](#)

In the Fact Editor, this fact is listed as `vbridge.health`:

```
<fact name="vbridge.health" value="true" type="Boolean" />
```

Attached Virtual NICs: This list box lists the Virtual NICs that are attached to this vBridge.

In the Fact Editor, this fact is listed as an array:

```
<fact name="vbridge.vnics">
  <array type="String">
    </array>
  </fact>
```

The list includes vNICs from currently running VMs only. When a VM is not running, PlateSpin Orchestrate does not consider its vNICs as “currently attached.” This is because vNICs are configured to associate with Network objects, and the attached vBridge might change, depending on which host the VM has been provisioned. The list can also change dynamically if the VM is migrated to another host (and so that host’s vBridge) on the same Network.

The Virtual Bridge Groups Panel

This section of the *Info/Groups* page lists the groups of vBridge objects (called “Networks”) in the grid to which this vBridge is associated. Click *Choose* to open the Network Selection dialog box. In this dialog box, you can choose which networks to display in the Explorer tree by selecting a group and then clicking *Add* or *Remove* to move it to or from the *Source Networks* list.

9.2.2 The Virtual Bridge Policies Tab

The *Policies* tab opens a page that contains a policy viewer for each of the policies associated with a Grid object. Click *Choose* to associate an existing policy with a vBridge.

NOTE: You can edit a policy by right-clicking a policy icon, selecting *Edit Policy* and clicking the Save icon.

9.2.3 The Virtual Bridge Health Debugger Tab

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Appendix A.3, “The Health Debugger,” on page 187](#).

9.2.4 The Virtual Bridge Constraints/Facts Tab

The Constraints/Facts tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. In essence, by building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any and all system resources by managing an object’s facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode *r/o* have read-only values, which can be viewed (that is, using the edit “pencil” icon) but changes cannot be made.

The Repository Object

10

Repositories are storage areas for VM image files and VM template files.

If a VM's files are stored on a particular host server, that VM must be run from that host server. If a VM's files are stored on a shared repository, that VM can be run on any host server that has access to the shared repository.

Host servers can have multiple repositories associated with them, and some repository types can be associated with multiple host servers as shared repositories. A host server can be associated with repositories stored locally on its server and with shared repositories stored on other machines.

The default size for all repositories is unlimited. To control disk space usage, you can change this default. For more information, see [Capacity \(MB\)](#): in “[Repository Information Subpanel](#)” on [page 126](#).

The repository groups and their constituent repository objects are displayed in the Explorer panel and the accompanying Repository Admin View of the Development Client.

- ♦ [Section 10.1, “Right-Click Menu Actions on the Repository Object,” on page 125](#)
- ♦ [Section 10.2, “Repository Groups,” on page 126](#)
- ♦ [Section 10.3, “The Repository Info/Groups Tab,” on page 126](#)
- ♦ [Section 10.4, “The Repository Policies Tab,” on page 132](#)
- ♦ [Section 10.5, “The Repository Health Debugger Tab,” on page 133](#)
- ♦ [Section 10.6, “The Repository Constraints/Facts Tab,” on page 133](#)
- ♦ [Section 10.7, “The Repository Action History Tab,” on page 133](#)
- ♦ [Section 10.8, “Repository Object Naming and Renaming,” on page 133](#)

10.1 Right-Click Menu Actions on the Repository Object

The Repository object displayed in the Explorer tree has three available actions in the right-click menu.

- ♦ **Discover VM Images:** Use this action on a newly-discovered or a newly-created repository to populate it with the [VMs](#) residing in the [mapped datastore](#) location.
- ♦ **Delete:** Use this action to delete the Repository object from the Explorer tree in the Orchestrate Development Client.

Before you delete the Repository object, you need to delete all the VMs contained in that repository; otherwise, an error message is displayed.

- ♦ **Rename:** Use this action to rename the Repository object. For more information, see [Section 10.8, “Repository Object Naming and Renaming,” on page 133](#).



10.2 Repository Groups

Any group object displayed in the Explorer panel represents a collection of similar object types. Groups can also be created automatically, such as when a provisioning adapter (PA) discovers a local repository on a VM host. For example, the Xen PA, upon discovery of a VM host, automatically creates a local repository for that VM host and places the created repository in a Xen repository group. You can also create groups manually in the Development Client, either by clicking the *Actions* menu and choosing *Create Repository Group* or by right-clicking the Repository object (anywhere in the Repository hierarchy) and selecting *New Repository Group*.

10.3 The Repository Info/Groups Tab

The page that opens under the *Info/Configuration* tab includes several collapsible sections on the page where you can configure the general information and attributes of the repository.

- ♦ [Section 10.3.1, “The Info Panel,” on page 126](#)
- ♦ [Section 10.3.2, “Best Practices for Entering Repository File Paths,” on page 131](#)
- ♦ [Section 10.3.3, “Groups,” on page 132](#)

NOTE: Whenever you make changes to any Grid object, the write icon  is superimposed on the object’s icon, signifying that the object has been altered. If you want to save the changes you have made, you need to click the Save icon  on the Development Client toolbar.

10.3.1 The Info Panel

The following fields on the Information panel provide facts for the Repository object:

- ♦ [“The “Show Inherited Fact Values” Check Box” on page 126](#)
- ♦ [“Repository Information Subpanel” on page 126](#)
- ♦ [“SAN Adapter Configuration” on page 131](#)

The “Show Inherited Fact Values” Check Box

Select this check box to show facts with overridden values supplied through attached or inherited policies. These fact values are read only (non-editable).

Repository Information Subpanel

The Repository Information panel on the Info/Groups page includes the following fields:

NOTE: Tooltip text is available when you mouse over any of these fields.

Description: The nature or purpose of this repository.

In the Fact Editor, this fact is listed as `repository.description`:

```
<fact name="repository.description" value="" type="String" />
```

Repository Enabled: This check box is selected by default. When it is selected (it has a value of true), VMs can be moved to this repository or they can be provisioned from it.

In the Fact Editor, this fact is listed as `repository.enabled`:

```
<fact name="repository.enabled" value="true" type="Boolean" />
```

Healthy: This check box is selected by default. When it is selected (it has a value of `true`), the repository is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy. For more information, see [Appendix A, “Grid Object Health Monitoring,” on page 185](#)

In the Fact Editor, this fact is listed as `repository.health`:

```
<fact name="repository.health" value="true" type="Boolean" />
```






Type: Select the repository type for this Repository object by selecting an option from the drop-down list.

In the Fact Editor, this fact is listed as `repository.type`:

```
<fact name="repository.type" value="local" type="String" />
```

The following table includes information about the various repository types:

Table 10-1 *Repository Types and Descriptions*

Repository Type	Description	Development Client Icon
local	The default repository on a host server. Each host server starts with its own local repository, which has the same name as the server's Resource Grid object.	
NAS (Network Attached Storage)	Represents a NAS device connected to host servers (for example, NFS mount). This NAS device must be mounted and available on all host servers associated with this Repository Grid object.	
SAN (Storage Area Network)	Represents an iSCSI or Fibre Channel SAN. Currently supported only with the vsphere provisioning adapter.	
datagrid	The shared datagrid repository (named <code>zos</code>) is located on the Orchestrate Server and is accessible to all host servers in the datagrid. By default, each host server has access to the <code>zos</code> datagrid repository.	
virtual	Represents an externally managed virtual disk (for example, Amazon EC2).	

IMPORTANT: If you have a vSphere environment with an iSCSI datastore based on an ESX 3.5 (or previous) host, the Development Client incorrectly displays its type as `local` rather than `SAN`. This misrepresentation affects the accuracy of the VM host assignment (how the repositories are scored in the plan), and could possibly affect VM migration validation.

To work around this issue, set the type to `SAN`. You need to check that this setting is retained when another VM host or Repository discovery is executed.

SAN ID: (SAN repositories only) The SAN ID (the Virtual Fabric ID) for this repository.

In the Fact Editor, this fact is listed as `repository.id`:

```
<fact name="repository.id" value="test1" type="String" />
```

Root Location: The repository's logical root location. You can also think of this as the base location for all VM files and subdirectories contained within this repository.

In the Fact Editor, this fact is listed as `repository.location`:

```
<fact name="repository.location" value="/" type="String" />
```

The table below provides some examples you can consider as you enter a shared root path in this field. For more information, see [“Best Practices for Entering Repository File Paths” on page 131](#).

Table 10-2 *Repository Types and Root Location Examples*


Repository Type	Root Location Examples
local	<ul style="list-style-type: none">♦ / (root)♦ c:/vm
NAS (Network Attached Storage)	This is the mount point that is assumed to be the same on every host server with a connection to this NAS. <ul style="list-style-type: none">♦ /u/mnt/myshareddisk
SAN (Storage Area Network)	Not required.
datagrid	grid:///vms
virtual	<ul style="list-style-type: none">♦ / (root)♦ c:/vm

VM Config Search Path: The relative path (from `repository.location`) to be used during discover of VM configuration files. This fact also implicitly includes the `resource.preferredpath` fact. For xen30 repositories, the default path is `/etc/xen/vm`.

In the Fact Editor, this fact is listed as an array:

```
<fact name="repository.searchpath">
  <array>
    <string>/etc/xen/vm</string>
  </array>
</fact>
```

IMPORTANT: If you use this field, do not include a leading forward slash (/) in the path. For more information, see [“Best Practices for Entering Repository File Paths” on page 131](#).

The  button opens the Attribute element values dialog box, where you can add, remove, or edit the path (element values) in an array of path choices.

The table below provides some examples you can consider as you enter a search path in this field.

Table 10-3 Repository Types and VM Config Search Path Examples

Repository Type	VM Config Search Path Examples
local	/etc/xen/vm
NAS (Network Attached Storage)	Each of these is the relative path from the location to search for VM configuration files. Null specifies to search the whole mount. <ul style="list-style-type: none">♦ <i>my_vms</i>♦ <i>saved_vms</i>♦ null (no path entry)
SAN (Storage Area Network)	Not required.
datagrid	grid:///vms
virtual	

Preferred Storage Path: The relative path (from `repository.location`) where you want PlateSpin Orchestrate to place the VM files after a move or a clone operation.

In the Fact Editor, this fact is listed as `repository.preferredpath`:

```
<fact name="repository.preferredpath" value="" type="String" />
```

IMPORTANT: If you use this field, do not include a leading forward slash (/) in the path. For more information, see [“Best Practices for Entering Repository File Paths” on page 131](#).

Table 10-4 Repository Types and Preferred Storage Path Examples

Repository Type	Preferred Storage Path Examples
local	var/lib/xen/images for Xen VMs
NAS (Network Attached Storage)	<i>my_vms</i>
SAN (Storage Area Network)	Not required.
datagrid	grid:///vms
virtual	

Capacity (MB): The maximum amount (in MB) of storage space available to VMs. The default (-1) designates an unlimited amount of space.

In the Fact Editor, this fact is listed as `repository.capacity`:

```
<fact name="repository.capacity" value="-1" type="Integer" />
```

Used Space (MB): The amount (in MB) of storage space used for VMs.

In the Fact Editor, this fact is listed as `repository.usedspace`:

```
<fact name="repository.usedspace" value="0" type="Integer" />
```

Free Space (MB): The amount (in MB) of storage space available to new VMs. The value is always set to -1, which designates an unlimited amount of space.

In the Fact Editor, this fact is listed as `repository.freespace`:

```
<fact name="repository.freespace" value="-1" type="Integer" />
```

Efficiency: Enter an efficiency coefficient that PlateSpin Orchestrate uses to calculate the cost of moving VM disk images to and from the repository. This value is multiplied by the disk image size (in MB) to determine an efficiency score. A score of zero (0) means no cost (very efficient).


In the Fact Editor, this fact is listed as `repository.efficiency`:

```
<fact name="repository.efficiency" value="1.0000" type="Real" />
```

Stored VMs: The VM images stored in this repository. The list is aggregated from individual VM facts.

In the Fact Editor, this fact is listed as an array:


```
<fact name="repository.vmimages">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the  button to open the *Attribute Element Values* dialog box, where you can add, remove, or edit the VM host IDs (element values) in an array of VM host ID choices.

Compatible VM Hosts: The VM hosts capable of using this repository. The list is aggregated from individual VM facts.

In the Fact Editor, this fact is listed as an array:

```
<fact name="repository.vmhosts">
  <array>
    <string>tszen4_xen30</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open the Attribute element values dialog box, where you can add, remove, or edit the VM host IDs (element values) in an array of VM host ID choices.

Accessed By Provision Adapters: The provisioning adapter jobs that are allowed access to VMs on this repository.

In the Fact Editor, this fact is listed as an array:

```
<fact name="repository.provisioner.jobs">
  <array>
    <string>xen30</string>
  </array>
</fact>
```

You can edit this array by clicking the  button to open the Choose Grid Objects dialog box, where you can add or remove provisioning adapters for the array of provisioning adapter choices.

NOTE: In the Fact Editor, you edit the provisioning adapter array by using the *Attribute Element Values* dialog box.

SAN Adapter Configuration

SAN Adapter Vendor: (SAN repositories only) The name of the vendor of the SAN. This should be adapter specific, such as *iqn*, *npiv*, *emc*. An empty field (that is, no value in the string) indicates that bind/unbind is a noop (no operation performed).

In the Fact Editor, this fact is listed as `repository.san.vendor`:

```
<fact name="repository.san.vendor" value="" type="String" />
```

SAN Transport: (SAN repositories only) From the drop-down list, select *iSCSI* or *Fibre Channel* to indicate the type of SAN transport this repository uses.

In the Fact Editor, this fact is listed as `repository.san.type`:

```
<fact name="repository.san.type" value="" type="String" />
```

10.3.2 Best Practices for Entering Repository File Paths

Use the following guidelines in scenarios where you need VM repositories.

- ♦ “Creating a Repository to Use with New VMs” on page 131
- ♦ “Creating a Repository to Use with Existing VMs” on page 131
- ♦ “Creating a Repository for Existing VMs with Shared Root Locations and Separate Configuration Directories” on page 132

Creating a Repository to Use with New VMs

If you are creating a repository for new VMs that you will eventually provision:

- 1 In the *Root Location* field, specify the location for the new repository.
Example: `/vms_new`
- 2 In the *Preferred Storage Path* field, specify the path to your image file store (relative to the root location path). This becomes the path for VM configuration files and VM image files when you associate a VM with this repository.
Example: `images` (no leading forward slash)
Because the fields are concatenated, the provisioning adapter searches for the existing VM files in `/vms_new/images`.

Creating a Repository to Use with Existing VMs

Use this procedure when you already have VMs in your grid and a store for the VM configuration and disk image files already exists.

- 1 In the *Root Location* field, specify the shared location for this repository.
Example: `/vms_new`
- 2 In the *VM Config Search Path* field, specify the search path to your existing configuration file store (relative to the Root Location path).

Example: `old_config` (no leading forward slash)

Because the fields are concatenated, the provisioning adapter searches for the existing VM configuration files in `/vms_new/old_config`.

- 3 In the *Preferred Storage Path* field, specify the path to your existing image file store (relative to the root location path). This also becomes the path for VM configuration files and VM image files when you associate a VM with this repository.

Example: `all_images` (no leading forward slash)

Because the fields are concatenated, the provisioning adapter searches for the existing VM files in `/vms_new/all_images`.


Creating a Repository for Existing VMs with Shared Root Locations and Separate Configuration Directories

Use this procedure when you want to create a repository for existing VMs that have a shared root path but separate configuration file directories such as `/vms_new/old_config1` and `/vms_new/old_config2`.

- 1 In the *Root Location* field, specify the shared location for this repository.

Example: `/vms_new`

- 2 In the *VM Config Search Path* field, specify the search paths to your existing configuration file store (relative to the Root Location path).

Example: Adjacent to the *VM Config Search Path* field, click , click *Add element*, enter `old_config1` (no leading forward slash), click *OK*, click *Add element* again, specify `old_config2` (no leading forward slash), then click *OK*.

Because the fields are concatenated, the provisioning adapter searches for the existing VM configuration files in the array consisting of `/vms_new/old_config1` and `/vms_new/old_config2`.

- 3 In the *Preferred Storage Path* field, specify the path to your existing image file store (relative to the Root Location path). This path also becomes the path for VM configuration files and VM image files after a move or clone when a VM has been associated with this repository.

Example: `all_images` (no leading forward slash)

Because the fields are concatenated, the provisioning adapter searches for the existing VM files in `/vms_new/all_images`.

10.3.3 Groups

This section of the Info/Groups page lists the groups of Repository objects in the grid. Click *Choose* to open the repository Group selection dialog box. In this dialog box, you can choose which Repository Groups to display in the Explorer Panel by selecting a group, then clicking *Add* or *Remove* to move it to or from the *Source Repository Groups* list.

10.4 The Repository Policies Tab

The *Policies* tab opens a page that contains a policy viewer for each of the policies associated with a Grid object.

NOTE: You can edit a policy by right-clicking a policy icon, selecting *Edit Policy*, then clicking the Save icon.

10.5 The Repository Health Debugger Tab

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Appendix A.3, “The Health Debugger,” on page 187](#).

10.6 The Repository Constraints/Facts Tab

The *Constraints/Facts* tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties.

By building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any system resources by managing an object’s facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode *r/o* have read-only values, which can be viewed by using the pencil icon, but changes cannot be made.

10.7 The Repository Action History Tab

The *Action History* tab is displayed in the administrative view of the Repository object. When you select the *Action History* tab, a table displays a list of the history for all actions performed on this Grid object.

The Orchestrate Server must be connected to an audit database for the *Include Audit Database* check box to be available. If the *Include Audit Database* check box is selected in this view, the action status is not polled. Click the refresh icon in the toolbar to retrieve and display fresh data.

For more details about the information listed on the Action History page, see “[Action History in Admin Views of the Development Client](#)” in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

10.8 Repository Object Naming and Renaming

Some resource names are generated by the Orchestrate system and can therefore receive generic, arbitrary names such as *host1*, *host2*, *host3*, and so on. Repositories you name at creation time might also change later in their purpose or facilities.

As the quantity of these Repository objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object’s “display name” is visible in the Development Client interface, the VM Client interface, the Server Portal, and in optional *zos* and *zosadmin* commands.

NOTE: Repository object groups (that is, the folders that contain these Repository objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A Repository object’s name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a Repository object in the Orchestrate Development Client using one of three methods:

- ♦ Right-click the Repository object in the Explorer tree, then select Rename to allow editing of the display name.
- ♦ Triple-click the Repository object in the Explorer tree to allow editing of the display name.
- ♦ In the *Constraints/Facts* page, where you can select the Repository object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you begin to rename using one of these methods, you will notice that the fact value is pre-populated with the `#{objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the Repository object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the `zos` or `zosadmin` command line, see the [PlateSpin Orchestrate 2.5 Command Line Reference](#).

A User object represents an individual account that is allowed to connect to the PlateSpin Orchestrate Server. Administrator users are also allowed to connect by using the `zosadmin` command line and the PlateSpin Orchestrate Development Client user interfaces.

You can use the Orchestrate Development Client user interface to manually create a User object. You can create objects automatically if authentication through LDAP or Active Directory is enabled, or optionally if autoregistration is configured.

The user object icon and the red square user object icon.

- ♦ [Section 11.1, “User Groups,” on page 135](#)
- ♦ [Section 11.2, “The User Info/Groups Tab,” on page 135](#)
- ♦ [Section 11.3, “The User Policies Tab,” on page 141](#)
- ♦ [Section 11.4, “The User Health Debugger Tab,” on page 141](#)
- ♦ [Section 11.5, “The User Constraints/Facts Tab,” on page 142](#)
- ♦ [Section 11.6, “The User Action History Tab,” on page 142](#)



11.1 User Groups

Any group object displayed in the Explorer panel represents a collection of similar object types. Groups can also be created automatically, such as when a provisioning adapter (PA) discovers a local repository on a VM host. For example, the Xen PA, upon discovery of a VM host, automatically creates a local repository for that VM host and places the created repository in a Xen repository group. You can also create groups manually in the Development Client, either by clicking the *Actions* menu and choosing *Create User Group* or by right clicking a User Group object (anywhere in the User hierarchy) and selecting *New User Group*.

11.2 The User Info/Groups Tab

The page that opens under the *Info/Configuration* tab of the User admin view includes several collapsible sections on the page where you can configure the general information and attributes of the user.

- ♦ [Section 11.2.1, “Info,” on page 135](#)
- ♦ [Section 11.2.2, “Groups,” on page 141](#)

NOTE: Whenever you make changes to any Grid object, the write icon  is superimposed on the object’s icon, signifying that the object has been altered. If you want to save the changes you have made, you need to click the save icon  on the Development Client toolbar.

11.2.1 Info

The following fields on the Information panel provide facts for the User object:

- ♦ [“Show Inherited Fact Values Check Box” on page 136](#)

- ♦ [“User Information” on page 136](#)
- ♦ [“Personal Information” on page 137](#)
- ♦ [“Job Information” on page 138](#)
- ♦ [“Accounting Information” on page 138](#)
- ♦ [“Job Control” on page 139](#)
- ♦ [“Quota Information” on page 140](#)

Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached or inherited policies. These fact values are read only (non-editable).

User Information

The User Information panel on the Info/Groups page includes the following fields:

NOTE: Tooltip text is available when you mouse over any of these fields.

Account Enabled: This check box is selected by default. When the check box is selected, the user is allowed to log in and run jobs..

In the Fact Editor, this fact is listed as `user.enabled`:

```
<fact name="user.enabled" value="true" type="Boolean" />
```

Online: When this check box is selected (it has a value of true), the user is currently logged in to the server.

In the Fact Editor, this fact is listed as `user.online`:

```
<fact name="user.online" value="false" type="Boolean" />
```

Healthy: This check box is selected by default. When it is selected (it has a value of true), the user is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy. For more information, see [Appendix A, “Grid Object Health Monitoring,” on page 185](#)


In the Fact Editor, this fact is listed as `user.health`:

```
<fact name="user.health" value="true" type="Boolean" />
```

External Groups: A list of external groups (for example, LDAP) that this user belongs to.

In the Fact Editor, this fact is listed as an array:

```
<fact name="user.external.groups">
  <array type="String">
    </array>
</fact>
```

You can edit this array by clicking the  button to open the Attribute element values dialog box, where you can add, remove, or edit the name and value for every user environment you want to use.

Personal Information

First Name: The user's first name.

In the Fact Editor, this fact is listed as `user.name.first`:

```
<fact name="user.name.first" value="" type="String" />
```

Last Name: The user's last name.

In the Fact Editor, this fact is listed as `user.name.last`:

```
<fact name="user.name.last" value="" type="String" />
```

Password: The user's hashed login password.

In the Fact Editor, this fact is listed as `user.password`:

```
<fact name="user.password" value="kNLj1_Fc96C3ajVXcqQEGZBRrbivgxhhzK3TKLpP" type="String" />
```

IMPORTANT: The System User password should not be changed. If you try to change this password by using the Orchestrate Development Client, the password does not actually change and no warning is posted that the change was unsuccessful. There is no indication of authentication failure until you try to log in with the “new” password.

Email: The user's e-mail address.

In the Fact Editor, this fact is listed as `user.name.email`:

```
<fact name="user.name.email" value="" type="String" />
```

City: The name of the city where the user is located.

In the Fact Editor, this fact is listed as `user.location.city`:

```
<<fact name="user.location.city" value="" type="String" />
```

State: The name of the state or province where the user is located.

In the Fact Editor, this fact is listed as `user.location.state`:

```
<fact name="user.location.state" value="" type="String" />
```

Country: The name of the country where the user is located.

In the Fact Editor, this fact is listed as `user.location.country`:

```
<fact name="user.location.country" value="" type="String" />
```

Site: The name of the site (for example, a campus or building) where the user works.

In the Fact Editor, this fact is listed as `user.location.site`:

```
<fact name="user.location.site" value="" type="String" />
```

Environment: A list of default user environment variable names and values that the Orchestrate Server sets when executing joblets remotely.

In the Fact Editor, this fact is listed as a dictionary:

```
<fact name="user.env">
  <dictionary>
    <dictelement key="dfadsafd">
      <string>safdaf</string>
    </dictelement>
  </dictionary>
</fact>
```

Job Information

Total Job Count: The total number of jobs that this user has historically initiated on this Orchestrate server.

In the Fact Editor, this fact is listed as `user.history.jobcount`:

```
<fact name="user.history.jobcount" value="0" type="Integer" />
```

Active Jobs: The number of top-level jobs run with this user account that are in an active state.

In the Fact Editor, this fact is listed as `user.jobs.active`:

```
<fact name="user.jobs.active" value="0" type="Integer" />
```

Queued Jobs: The number of top-level jobs run with this user account that are currently in a queued state.

In the Fact Editor, this fact is listed as `user.jobs.queued`:

```
<fact name="user.jobs.queued" value="0" type="Integer" />
```

Total Jobs: The total number of top-level jobs run by this user account.

In the Fact Editor, this fact is listed as `user.jobs.total`:

```
<fact name="user.jobs.total" value="0" type="Integer" />
```

Active Sessions: The number of currently active sessions (that is, connections) that the user has established with the Orchestrate Server.

In the Fact Editor, this fact is listed as `user.sessions`:

```
<fact name="user.sessions" value="1" type="Integer" />
```

Accounting Information

Total Spending: The total cost of computing resources by this user.

In the Fact Editor, this fact is listed as `user.history.cost.total`:

```
<fact name="user.history.cost.total" value="0.0088" type="Real" />
```

Average Spending Rate: The computed moving average spending (in dollars per hour) over the last hour of activity for this user.

In the Fact Editor, this fact is listed as `user.account.spendrate`:

```
<fact name="user.account.spendrate" value="-0.0006" type="Real" />
```

Maximum Spending Rate: An amount (in dollars per hour) to be used by the Resource Scheduler to throttle the rate at which computing cycles are consumed by the user. A value of less than or equal to zero (≤ 0) turns the feature off.

In the Fact Editor, this fact is listed as `user.account.maxspendrate`:

```
<fact name="user.account.maxspendrate" value="0.0000" type="Real" />
```

Default Accounting Group: Lets you select the default User Group to be billed for work conducted by this user.

In the Fact Editor, this fact is listed as `user.accounting.group`:

```
<fact name="user.accountinggroup" value="all" type="String" />
```

Total Wall Time: The total amount of wall time (in seconds) consumed by this user.

In the Fact Editor, this fact is listed as `user.history.time.total`:

```
<fact name="user.history.time.total" value="31" type="Integer" />
```

Total Grid Time: The total amount of grid time (in gcycles, which is a normalized average of compute cycles) consumed by this user.

In the Fact Editor, this fact is listed as `user.history.gcycles.total`:

```
<fact name="user.history.gcycles.total" value="31" type="Integer" />
```

Job Control

Default Priority Value: A numerical representation of the default priority at which this user's job runs, with 1 being the lowest priority and 9 being the highest priority.

In the Fact Editor, this fact is listed as `user.priority.default`:

```
<fact name="user.priority.default" value="7" type="Integer" />
```

Default Priority: The string representation of the default priority at which this user can run a job. The value is matched to the integer value in `user.priority.default`.

In the Fact Editor, this fact is listed as `user.priority.default.string`:

```
<fact name="user.priority.default.string" value="high" type="String" />
```

Maximum Priority Value: A numerical representation of the maximum priority at which this user's job can run, with 1 being the lowest priority and 9 being the highest priority. Only the system user can run jobs at priority 10.

In the Fact Editor, this fact is listed as `user.priority.max`:

```
<fact name="user.priority.max" value="5" type="Integer" />
```

Datagrid Maximum History: The maximum number of job instance directories that should be kept in the datagrid for this user.

In the Fact Editor, this fact is listed as `user.datagrid.maxhistory`:

```
<fact name="user.datagrid.maxhistory" value="25" type="Integer" />
```

Job Preemption Enabled: Select this check box if you want to allow the user to preempt jobs that have a priority less than the priority of the running job instance.

In the Fact Editor, this fact is listed as `user.preemption.enabled`:

```
<fact name="user.preemption.enabled" value="false" type="Boolean" />
```

Max Preemption Priority: The highest job priority band from which this user is allowed to preempt resources. The value acts as a delta from the current job instance priority. The maximum preemptible priority is always less than or equal to `user.priority.max`.

In the Fact Editor, this fact is listed as `user.preemption.priority.delta`:

```
<fact name="user.preemption.priority.delta" value="0" type="Integer" />
```

Resources Stealing Enabled: Select this check box to allow the user to steal resources that are running jobs that have a priority less than the priority of the running job instance.

In the Fact Editor, this fact is listed as `user.stealing.enabled`:

```
<fact name="user.stealing.enabled" value="false" type="Boolean" />
```

Max Stealing Priority: The highest job priority band from which this user is allowed to steal resources. The value acts as a delta from the current job instance priority, and must be less than zero (<0).


In the Fact Editor, this fact is listed as `user.stealing.priority.delta`:

```
<fact name="user.stealing.priority.delta" value="-1" type="Integer" />
```

Privileged Job Groups: A list of Job Groups with jobs and joblets that this user is allowed to run on resources that have reached their slot maximum or that are provisioned resources that are reserved for another user or job.

In the Fact Editor, this fact is listed as an array:

```
<fact name="user.privilegedjobgroups">
  <array type="String">
    </array>
</fact>
```

You can edit this array by clicking the  button to open the Choose Grid Objects dialog box, where you can add or remove Job Groups in an array of choices. The Job Groups can be added to or removed from a list of Source Grid Objects to a list of Target Grid Objects (or vice versa).

Quota Information

Account Balance Remaining: The balance (measured in dollars) that remains available for this user since the last reset. You can use this value to implement quotas on your server.

In the Fact Editor, this fact is listed as `user.account.balance`:

```
<fact name="user.account.balance" value="0.0000" type="Real" />
```

Job Counter: The number of jobs this user has initiated since the last reset. You can use this value to implement quotas on your server.

In the Fact Editor, this fact is listed as `user.jobcount`:

```
<fact name="job.history.jobcount.complete" value="0" type="Integer" />
```

Time Remaining: The amount of wall time (measured in seconds) remaining for use by this user since last the reset. You can use this value to implement quotas on your server.

In the Fact Editor, this fact is listed as `user.account.time`:

```
<fact name="user.account.time" value="0" type="Integer" />
```

Grid Time Remaining: The amount of grid time (measured in gcycles) remaining for use by this user since last the reset. You can use this value to implement quotas in your grid.

In the Fact Editor, this fact is listed as `job.history.jobcount.complete`:

```
<fact name="user.account.gcycles" value="0" type="Integer" />
```

11.2.2 Groups

This section of the Info/Groups page lists the groups of User objects in the grid. Click *Choose* to open the User Group Selection dialog box. In this dialog box, you can choose which User Groups to display in the Explorer Panel by selecting a group and then clicking *Add* or *Remove* to move it to or from the *Source Job Groups* list.

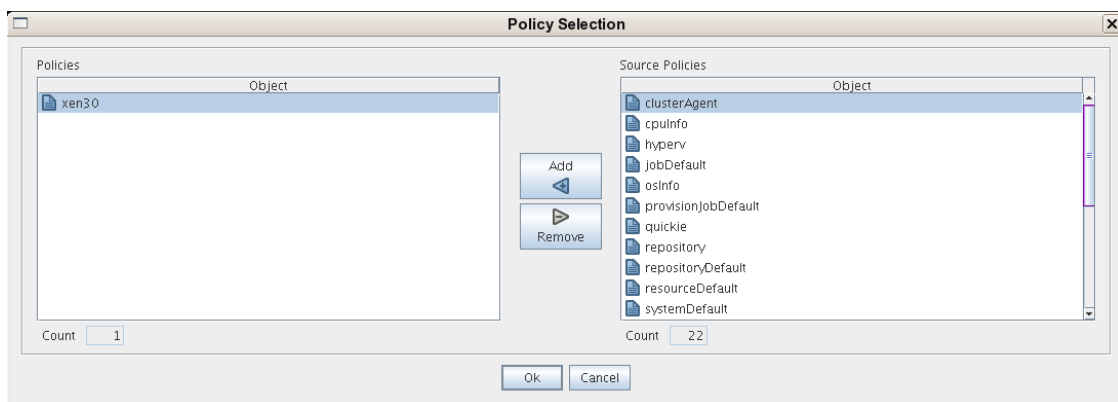
11.3 The User Policies Tab

The *Policies* tab of the User admin view opens a page that contains a policy viewer for each of the policies associated with a User Grid object.

You can modify policies by using the Policy Grid object. For more information, see [Section 12.1, “The Policy Object,” on page 143](#).

If you click *Choose* on the *Policy* tab, a Policy Selection dialog box is launched where you can add or remove individual policies to be applied to the selected User Grid object.

Figure 11-1 The Policy Selection Dialog Box



11.4 The User Health Debugger Tab

The Health Debugger is a common Admin view in the Development Client for most Grid objects. For information about this tool, see [Chapter A.3, “The Health Debugger,” on page 187](#).

11.5 The User Constraints/Facts Tab

The *Constraints/Facts* tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties.

By building, deploying, and running jobs on the PlateSpin Orchestrate Server, you can individually change the functionality of any system resources by managing an object's facts and constraints. The Orchestrate Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode *r/o* have read-only values, which can be viewed (that is, using the pencil icon), but changes cannot be made.

11.6 The User Action History Tab

The Action History tab is displayed in the administrative view of the User object. When you select the Action History tab, a table displays a list of the history for provisioning actions performed on this Grid object (assuming that it is provisionable, for example, a VM or VM template)..

The Orchestrate Server must be connected to an audit database for the *Include Audit Database* check box to be available. If the *Include Audit Database* check box is selected in this view, the action status is not polled. Click the refresh icon in the toolbar to retrieve and display fresh data.

For more details about the information listed on the Action History page, see “[Action History in Admin Views of the Development Client](#)” in the *PlateSpin Orchestrate 2.5 Virtual Machine Management Guide*.

Miscellaneous Objects Displayed in the Explorer Tree

12

The Explorer panel (tree) of the PlateSpin Orchestrate Development Client Explorer displays some miscellaneous objects of importance:

- ♦ [Section 12.1, “The Policy Object,” on page 143](#)
- ♦ [Section 12.2, “Computed Fact Objects,” on page 143](#)
- ♦ [Section 12.3, “Event Objects,” on page 144](#)
- ♦ [Section 12.4, “Metrics Objects,” on page 144](#)

12.1 The Policy Object

XML is used to define PlateSpin Orchestrate policies. A policy can be deployed to the server and associated with any grid object. The `policy` element is the root element for policies. Policies contain constraints and fact definitions for grid objects.

You can edit a policy by clicking its icon in the Explorer tree view, making your changes, then clicking the save icon. The *Where Used* tab of the Policy Editor lists the jobs where the selected policy is associated.

12.1.1 Policy Constraints

A policy can define a collection of constraints which are applied appropriately based on context. For example, a resource constraint can limit the selection of a resource to a subset based on resource group membership, or any number of other fact-based evaluations.

You can use `/opt/novell/zenworks/zos/server/examples/customNode.policy` on the PlateSpin Orchestrate Server as an example policy file with constraints.

12.1.2 Policy Facts

The XML fact element defines a fact to be stored in the grid object’s fact namespace. The name, type and value of the fact are specified as attributes. For list or array fact types, the element tag defines list or array members. For dictionary fact types, the `dict` tag defines dictionary members.

You can see an example policy with an XML representation for all the fact types on the PlateSpin Orchestrate Server at `/opt/novell/zenworks/zos/server/examples/allTypes.policy`.

12.2 Computed Fact Objects

Computed facts are used when you want to run JDL to generate the value for a fact. Although computed facts are not jobs, they use the same JDL syntax. You can see examples of computed facts on the PlateSpin Orchestrate Server at `/opt/novell/zenworks/zos/server/examples/activejobs.cfact` and `/opt/novell/zenworks/zos/server/examples/repostitory.cfact`.

12.3 Event Objects

An Event object in the Explorer tree represents a user-described set of rules that can be associated with a schedule trigger or handled by long-running jobs written to respond to events.

For more information about using events, see [“Using an Event Notification in a Job”](#) in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*, also see [“Event Triggers”](#) on page 153 of this guide, and [Section A.2, “Health Events,”](#) on page 187 of this guide.

12.4 Metrics Objects

The Orchestrate Metrics Facility collects, aggregates, and allows simple fact-based retrieval of metric values by jobs and computed facts (via JDL), policy constraints, and Event triggers on a per-resource basis.

A Metrics object is deployed in the Explorer tree. Use the right-click menu to display the “deploy” and “undeploy” actions. Pre-defined .metric files are located in the `/opt/novell/zenworks/zos/server/components/metrics` folder, or you can create a new .metric file and paste in an XML file.

NOTE: You can also use the `zosadmin deploy` command to deploy a .metric file. For example:

```
zosadmin deploy load_one.metric
```

Metrics objects are listed by their deployment name, which may or may not be the same as the name of the actual metric. This potentially allows multiple, separately deployable, RRD definitions for a single “instantaneous” metric, with different aggregation periods defined.

For more information about using metrics, see [Appendix C, “The Metrics Facility,”](#) on page 205 of this guide, and [“Resource Metrics Facts”](#) in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*.

The PlateSpin Orchestrate Job Scheduler

13

You can use the Job Scheduler in PlateSpin Orchestrate Server to automatically start deployed jobs on your grid by using either time or event triggers.

You can think of the functionality provided by the time triggers as being similar to a distributed cron system (in fact, time triggers can be described in cron syntax). This triggering, coupled with the job control functions in PlateSpin Orchestrate, allows for the sophisticated automation of routine data center tasks.

For example, suppose you want to periodically harvest a large log file in a coordinated way from a farm of several hundred machines. First, you could create an PlateSpin Orchestrate job that uses the datagrid for file movement. The job control options specify that the job should run on not more than three machines at once and sweep across the entire grid. You would then create a schedule to run this job at the desired interval.

As another example, you could use the Job Scheduler to trigger a discovery job every time a new resource is added to the grid. In this case, the job developer writes the discovery job to discover and set facts about the resource. Next, you would create a schedule to run this job on the `RESOURCE_ONLINE` built-in trigger. In fact, this type of triggered job is currently used in the standard set of deployed discovery jobs to detect specific resource CPU and OS information.

Yet another example would be to run a job on server startup that sends a notification e-mail to an administrator.

This section includes the following information:

- ♦ [Section 13.1, “Understanding the Job Scheduler View,” on page 145](#)
- ♦ [Section 13.2, “Walkthrough: Scheduling a System Job,” on page 162](#)

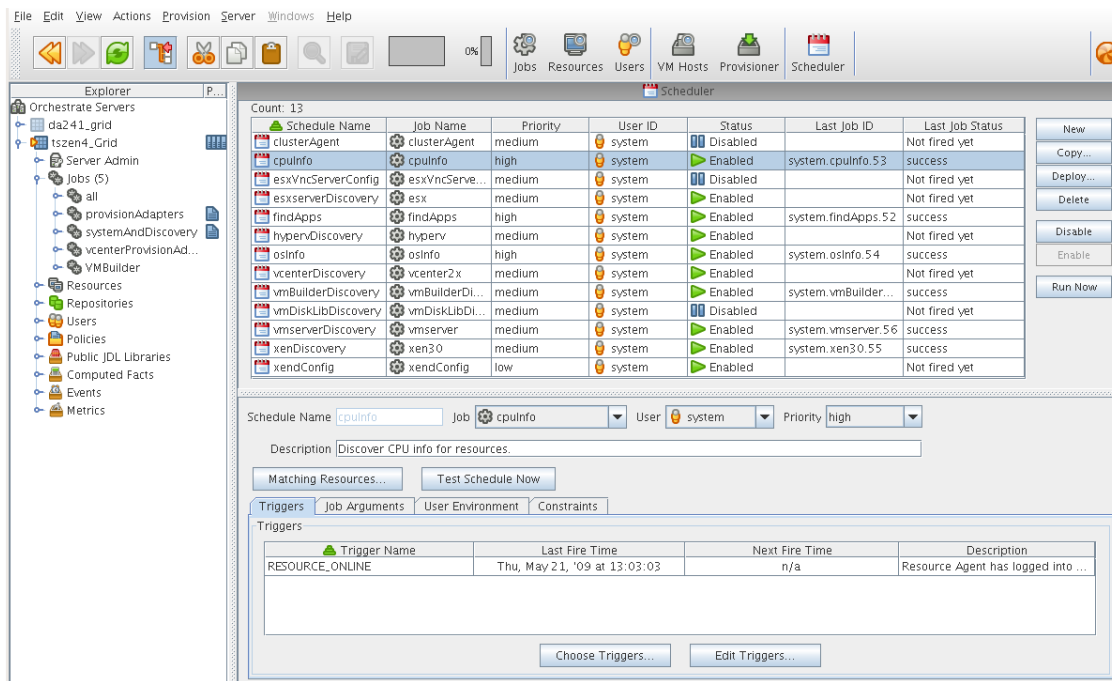
13.1 Understanding the Job Scheduler View

This section includes information to help you understand the functions of the Job Scheduler and how to use it to launch PlateSpin Orchestrate jobs.

- ♦ [Section 13.1.1, “Navigating The Job Schedules Table,” on page 146](#)
- ♦ [Section 13.1.2, “Creating or Modifying a Job Schedule,” on page 148](#)
- ♦ [Section 13.1.3, “Understanding Cron Syntax in the Job Scheduler,” on page 158](#)

Click *Scheduler* on the main toolbar of the PlateSpin Orchestrate Development Client to open the Job Scheduler view.

Figure 13-1 Job Scheduler View of the Orchestrate Development Client



13.1.1 Navigating The Job Schedules Table

PlateSpin Orchestrate includes several predefined and predeployed discovery jobs that have predefined launch schedules. Among these jobs are the `cpuinfo`, `findapps`, `osinfo`, and other jobs, depending on the options (that is, the “server profile”) you chose and the configuration you used during the installation. After installation, these jobs are listed by name in a table in the Job Scheduler view.

Figure 13-2 The Job Schedules Table in the Job Scheduler View

Schedule Name	Job Name	Priority	User ID	Status	Last Job ID	Last Job Status
cpuinfo	cpuinfo	high	zosSystem	Enabled	zosSystem.cpuInfo.2	success
findApps	findApps	high	zosSystem	Enabled	zosSystem.findAp...	success
osInfo	osInfo	high	zosSystem	Enabled	zosSystem.osInfo.3	success
vcenterDiscovery	vcenterDiscov...	medium	zosSystem	Enabled	zosSystem.vcenter...	success
vmBuilderDiscovery	vmBuilderDis...	medium	zosSystem	Enabled		Not fired yet
vmHostVncConfig	vmHostVncCo...	low	zosSystem	Disabled		Not fired yet
vmserverDiscovery	vmserverDisc...	medium	zosSystem	Enabled	zosSystem.vmserv...	success
whdDiscovery	whdDiscovery	medium	zosSystem	Enabled		Not fired yet
xenDiscovery	xenDiscovery	medium	zosSystem	Enabled	zosSystem.xenDis...	success

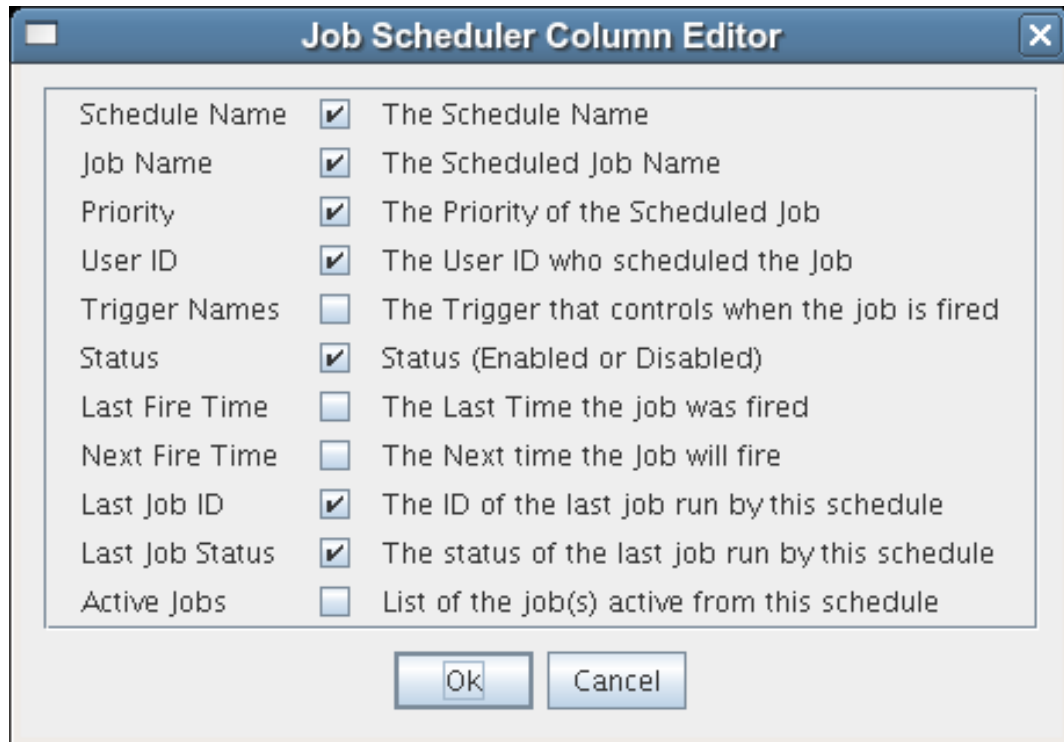
By default, PlateSpin Orchestrate uses schedule names that are similar to the job name so that schedules are easy to match (although this is not required). The schedules list shows all of the existing job schedules that accompany predefined jobs, along with the schedules that you create in the Job Scheduler.

NOTE: The Job Scheduler view is not a real-time monitor view of jobs, so if a job attribute (for example, Last Job Status or Last Fire Time) has changed, it might not be displayed until you click *Refresh*.

The Job Schedules Table has functionality that lets you decide how you want to display information about the job schedules:

- ♦ You can drag any column in the table to move it left or right in the table according to your preference.
- ♦ You can mouse over any column heading in the table to view tool tip text about the purpose of the data in that column.
- ♦ You can right-click any column heading in the table to open the Job Scheduler Column Editor dialog box.

Figure 13-3 Job Scheduler Column Editor Dialog Box



You can select any column heading in this dialog box to display it in the Job Schedules Table. The columns display the attributes of a previously configured job schedule. As the figure shows, this dialog box also includes text that clarifies the purpose of the data in each column.

In the Job Scheduler view, there are seven function buttons next to the Job Schedules Table (see [Figure 13-2 on page 146](#)) that let you take action on any schedule you select inside the table. (Only one schedule at a time can be selected.)

- ♦ **New:** Opens a dialog box where you can create a new schedule. When you create a new schedule, the Job Scheduler adds a new line to the Job Schedules Table. When the new line is added, you can use the Job Schedule Editor to edit the attributes of the schedule. A new schedule must be given a unique schedule name.

The Job Scheduler forces a new schedule to be created in the *Disabled* state to prevent it from running while it is being defined. You click *Enable* when a job is ready to be used.

- ♦ **Copy:** Copies a schedule you have selected in the Job Schedules Table. Clicking this button opens a dialog box where you rename the copy. If you want to create a schedule similar but not identical to an existing schedule, use this button to save time in adding attributes to a job schedule configuration. A copy of a schedule must be given a unique schedule name.
- ♦ **Deploy:** Opens a dialog box where you can select a schedule (that is, a deployable .sched file) to deploy.
- ♦ **Delete:** Deletes the selected schedule from the Job Schedules Table. You cannot recover a deleted job schedule.


























NOTE: Deleting a schedule that was deployed as part of a .job or .sar displays a confirmation dialog box. Deleting the schedule undeploys all contents of the .job or .sar that contains the schedule.

- ♦ **Disable:** Disables the selected schedule in the Job Schedules Table. The jobs associated with the schedule are not re-run, but any currently running instances of this job continue to run.
- ♦ **Enable:** Enables a disabled job schedule.
- ♦ **Run Now:** Forces the specified schedule to run immediately. This updates statistics such as *Last Fire Time*.

Removed Jobs or Users: Scheduler Behavior

If a job or a user is undeployed or removed from PlateSpin Orchestrate, the Job Schedules Table continues to list the schedule previously associated to that removed grid object, but the removed grid object no longer displays the icon that represents the object (job or user).

Figure 13-4 Some User Object and Job Object Icons Not Displayed

 Schedule Name	Job Name	Priority	User ID	Status	Last Job ID	Last Job Status
 CpuDiscovery	cpuInfo	high	 zosSystem	 Disabled		Not fired yet
 clusterAgent	quickle	medium	 zosSystem	 Disabled		Not fired yet
 findApps	findApps	high	 zosSystem	 Enabled	zosSystem.findAp...	success
 osInfo	osInfo	medium	 system	 Enabled	zosSystem.osInfo.48	success
 vcenterDiscovery	vcenterDiscov...	medium	 zosSystem	 Enabled	zosSystem.vcenter...	success
 vmHostVncConfig	vmHostVncCo...	low	 zosSystem	 Disabled		Not fired yet
 vmserverDiscovery	vmserverDisc...	medium	 zosSystem	 Enabled	zosSystem.vmserv...	success
 xenDiscovery	xenDiscovery	medium	 zosSystem	 Enabled	zosSystem.xenDis...	success

In the preceding figure, the *CpuDiscovery* schedule displays no Job icon for the *cpuInfo* job in the schedules table. Even though the job has been undeployed, the schedule is still listed.

In the *osinfo* schedule, the *system* user has no User icon. That user has been removed from PlateSpin Orchestrate.

If you choose a new user or job to be associated with a schedule, a deleted or undeployed user or job is never displayed in the popup menu for that schedule again.

13.1.2 Creating or Modifying a Job Schedule

The Job Schedule Editor is located immediately below the Job Schedules Table in the Job Scheduler view.

Figure 13-5 The Job Schedule Editor in the Job Scheduler View

Schedule Name: Job: User: Priority:

Description:

Triggers | Job Arguments | User Environment | Constraints

Triggers

Trigger Name	Last Fire Time	Next Fire Time	Description
RESOURCE_ONLINE	Thu, Jul 24, '08 at 16:43:23	n/a	Resource Agent has logged into server

There are several times when you can use this part of the Job Scheduler tool:

- ♦ When you create a new schedule by clicking *New*.
- ♦ When you modify the attributes of an existing schedule (available when you select a schedule in the table).
- ♦ When you create a copy of an existing schedule by clicking *Copy*.

The Job Schedule Editor lets you create or modify a job schedule by specifying its attributes.

You can use the following controls and data when you create or modify a job schedule:

- ♦ “Schedule Name” on page 149
- ♦ “Job” on page 149
- ♦ “User” on page 150
- ♦ “Priority” on page 150
- ♦ “Description” on page 150
- ♦ “Matching Resources” on page 150
- ♦ “Test Schedule Now” on page 150
- ♦ “Triggers” on page 150
- ♦ “Job Arguments” on page 156
- ♦ “User Environment” on page 157
- ♦ “Constraints” on page 158

Schedule Name

When you create a new schedule, the unique name you specify is displayed in this field. If you select a schedule from the Job Schedules Table, the name of the schedule is displayed in this field. The field is not editable, because schedules cannot be renamed after they are created. (You can use a copy if this is required.)

Job

When you create a new schedule, you need to associate a deployed job with it. You can select the job you want to run from this drop-down list.

If you want to use a previously created schedule to run a different job, you can change the job here.

User

When you create a new schedule, you need to associate a user with it. The user represents the user for whom the job will run. The choice of user might affect the permissions, privileges and constraints of the job. You can select the user from this drop-down list.

If you want a different user to run a job on a previously created schedule, you can change the user here.

If you decide to change the user who runs the job, check the *Priority* field to make sure that the priority you want is selected.

Priority

When you create a new schedule and associate a job and a user with it, a list of possible run priorities becomes available in this drop-down list. The list of priorities varies, depending on the user that is specified in the previous field. In this field, you select the priority of the job that is to be run so that if other jobs are to start concurrently or are competing for resources, PlateSpin Orchestrate can determine which job takes priority.

Description

For predeployed jobs, this field contains a default description of what the job's schedule does. The field is editable, so you can enter a description of your own for job schedules that you create.

Matching Resources

This button displays a list of resources where the job runs now or where it could run. This list is useful for checking the context of constraints that might have been affected by a choice of user or by manually specifying additional constraints under the *Policy* tab. The list is also useful to verify that a discovery job (that is, one that is triggered by the *Run on Resource Start* option) runs on the preferred set of machines.

Test Schedule Now

Click this button to test the new or modified schedule you are working with. The test runs the new or modified schedule without permanently saving the current configuration of the schedule or recording statistics. This control differs from the *Run Now* control in the Job Schedules Table, which runs a saved (persisted) schedule, disregarding any unsaved modifications that have been made to it in the Job Schedule Editor.

Triggers

When you click the *Triggers* tab in the Job Scheduler view, the following page opens:

Figure 13-6 The Schedule Triggers Page in the Job Scheduler

Schedule Name: Job: User: Priority:

Description:

Triggers | Job Arguments | User Environment | Constraints

Trigger Name	Last Fire Time	Next Fire Time	Description
RESOURCE_ONLINE	Thu, Jul 24, '08 at 16:43:23	n/a	Resource Agent has logged into server

In this view, you can add or define the triggers you want to associate with job schedules. A trigger is the signal to the Job Scheduler to initiate, or “fire” a schedule at a given time or at the occurrence of a given event. Job Scheduler triggers can be classified with regard to two conditions: events and time.

The Triggers table on this page has functionality that lets you decide how you want to display information about the triggers:

- You can drag any column in the table to move it left or right in the table according to your preference.
- You can mouse over any column heading in the table to view tool tip text about the purpose of the data in that column.
- You can right-click any column heading in the table to open the Triggers table column editor dialog box.

Figure 13-7 Trigger Table Column Editor Dialog Box

Trigger Name ☒ The Name of the Trigger

Last Fire Time ☒ The Last Time the job was fired

Next Fire Time ☒ The Next time the Job will fire

Description ☒ Trigger Description


You can select any column heading in this dialog box to display it in the Triggers table. The columns display the attributes of a previously configured Triggers table. As the figure shows, this dialog box also includes text that clarifies the purpose of the data in each column.

You can create as many triggers as you want to meet any scheduling situation you might have. Multiple time triggers can be associated with a schedule and multiple schedules can use the same trigger. The triggers you create are retained by the Job Scheduler for you to choose from when you create a schedule for a job. The currently associated triggers are displayed in the list along with a description.

Choose Triggers

This button opens a dialog box where you can choose both predefined and user defined time triggers to associate with this job schedule.

Figure 13-8 Choose Triggers Dialog Box



In this dialog box, you can click *Add* to move a selected trigger to the active, scheduled triggers that are to be associated with this job schedule. You can also click *Remove* to unassociate a trigger.

When a trigger is moved to the scheduled list, it becomes associated to the job schedule and it is displayed in the Job Scheduler view.

Most example jobs in PlateSpin Orchestrate are associated with event triggers, which are shown in the previous illustration. The dialog box can also list other job schedule triggers that are based on time.

Event Triggers

An Event trigger is the signal to the Job Scheduler to initiate, or “fire” a job when a given event occurs. An Event can be one of three types:

- ♦ **Event objects:** These objects are user defined events that are fired when an event rule is triggered. If an event object is deployed, it automatically shows in the trigger chooser as a possible choice.
- ♦ **built-in events:** These events are system wide events such as when a resource comes online or when a resource health condition change occurs. Built-in events are always available as a trigger choice. The Job Scheduler has eight possible built-In event triggers:
 - ♦ AGENT_VERSION_MISMATCH
 - ♦ RESOURCE_ONLINE
 - ♦ REPOSITORY_HEALTH
 - ♦ RESOURCE_HEALTH
 - ♦ SERVER_UP
 - ♦ USER_HEALTH
 - ♦ USER_ONLINE
 - ♦ VMHOST_HEALTH

You can select any combination of these event triggers for a single schedule.

The first trigger, AGENT_VERSION_MISMATCH, triggers the job when a PlateSpin Orchestrate Agent of an incompatible version attempts to connect to this Orchestrate Server. It can be used to initiate a configuration management tool for an agent software update or the job could e-mail an administrator to report the incompatible agent. The other seven available built-in event triggers are listed with accompanying descriptions in the dialog box.

- ♦ **External events:** These events are fired by an outside process. These are not automatically shown as choices in the trigger chooser, but must be defined by the trigger editor.

An event trigger can be used in conjunction with a time trigger to allow flexibility in scheduling the job application for maximum effectiveness or convenience. Jobs triggered by events require that their job arguments contain a dictionary named `context`. For example, your event-triggered job should have this jobarg element in its policy:

```
<policy>
  <jobargs>
    <fact name="context" type="Dictionary"
      description="Dictionary containing the context for the event" />
  </jobargs>
</policy>
```

The key/values of the dictionary are dependent on the event type. For event objects, the `jobargs.context` dictionary contains the matching context of the triggered rule. For built-in events, the `jobargs.context` dictionary contains the key of the object type corresponding to the built-in event and the object ID that caused the event.

For example, if the USER_ONLINE event triggers because the user named `foo` logs in, the `jobargs.context` dictionary contains:

```
{ user : foo }
```

Likewise, if the RESOURCE_ONLINE event is triggered because the resource agent named "vmhost1" comes online, the jobargs.context dictionary contains:

```
{ resource : vmhost1 }
```

For the AGENT_VERSION_MISMATCH event, the jobargs.context dictionary contains more information, as shown in the following table:

Table 13-1 Dictionary Information

Key	Type
AgentBuild	Long
AgentIP	String
AgentId	String
AgentMajor	Integer
AgentMinor	Integer
AgentPoint	Integer
JavaMajor	Integer
JavaMinor	Integer
JavaPoint	Integer
JavaVendor	String
JavaVersion	String
OsMajor	Integer
OsMinor	Integer
OsName	String
OsPoint	Integer
OsVendor	String
OsVersion	String
SystemArch	String
UsingJRE	Boolean
resource	String

Time Triggers

A time trigger is the signal to the Job Scheduler to initiate, or “fire” a job when a prescheduled time occurs. A time trigger can be used in conjunction with an event trigger to allow flexibility in scheduling the job application for maximum effectiveness or convenience. No default time triggers are defined in the Job Scheduler. You need to create new time triggers by clicking *Edit Triggers*.

Edit Triggers

Click *Edit Triggers* to open the Triggers dialog box.

Figure 13-9 The Triggers Dialog Box

Count: 0

Trigger Name	Trigger Description
--------------	---------------------

Trigger Name Description ...

☐ Fire Starting in Min Repeating every Min Repeat indefinitely

☐ Fire using CRON Expression
Examples

☒ Fire using Event ...
Events


Save

Close

The following controls and information are available in the dialog box:

- ♦ **New:** Opens a secondary dialog box where you can create a new time trigger name. When you create the trigger name, the attribute fields in the Triggers dialog box are cleared and you can specify new attributes for the trigger. A new trigger must be given a unique trigger name.
- ♦ **Copy:** Lets you modify an existing time trigger by giving it a new name and attributes. This can be helpful if there are only slight differences in the new attributes. A copy of a trigger must be given a unique trigger name.
- ♦ **Deploy:** Opens a file chooser where you can choose an existing, stored trigger (that is, a .trig file) to deploy.
- ♦ **Delete:** Deletes a selected time trigger.


IMPORTANT: Deleted triggers are not recoverable. If the trigger is used by existing schedules, it is removed from all of those schedules when it is deleted.

- ♦ **Trigger Name:** Specifies the unique name of the trigger you are creating or modifying. This name is displayed in the Job Scheduler if you choose to associate this trigger with a schedule. After you create the trigger name, it cannot be modified.
- ♦ **Description:** Specifies a description for the time trigger you are creating or modifying. The description is optional and can be as detailed as you want.
If the number of characters in the description string exceeds the space in the *Description* field, a  button is enabled that opens a string editor when clicked.
- ♦ **Save:** Clicking this icon saves the defined time trigger and its attributes.
- ♦ **Fire Starting In:** Displays multiple fields specifying the time increment and frequency to be used by the trigger to fire the job. If you select this type of time trigger, the *Fire using CRON Expression* button becomes inactive.

NOTE: You can use the *Fire Starting In* control to create either a “one-shot” time trigger or a “reoccurring” time trigger.

A one-shot time trigger fires just once after a specified period of time. To specify a one-shot trigger, click *Fire Starting in*, specify the amount of time before firing, then specify 0 as the time to *Repeat Indefinitely*.

A reoccurring time trigger fires after a specified period and then either fires repeatedly for an indefinite number of times or it fires for a specified number of times. To specify a reoccurring, indefinite trigger, click *Fire Starting in*, specify the amount of time before firing, then select *Repeat Indefinitely*. To specify a reoccurring but finite trigger, click *Fire Starting in*, specify the amount of time before firing, select *Repeat Range*, then specify the number of times you want the trigger to fire.

-
- ♦ **Fire using CRON Expression:** Specifies the cron expression that enables the job to fire automatically at a specified time or date. You need to be familiar with cron to use this field.
The *Examples* list box of selected cron expressions and their associated descriptions is located just below this button. You can use a listed expression as is, or use it as a template to modify the expression to meet your needs.
If you select this type of time trigger, the *Fire Starting In* and the *Fire Using Event* buttons become inactive.
For an example of how a cron expression can be implemented in a trigger, see [“Creating and Assigning a Time Trigger for the New Schedule” on page 167](#). For detailed information about cron syntax, see [“Understanding Cron Syntax in the Job Scheduler” on page 158](#).
 - ♦ **Fire Using Event:** Specifies a deployed event or an external event that enables the job to fire when a specified event occurs. Deployed events are defined using an XML syntax. You can specify a deployed event from *Events* (that is, listed in the *Events* drop down list) or you can enter the name of an external event. For more information on creating and firing an event, see
If the number of characters in the *Fire Using Event* description string exceeds the space in the field, a  button is enabled that opens a string editor when clicked.

Job Arguments

This tab displays an area (in the lower left corner of the Job Schedule Editor) where possible job arguments are listed. If you select an existing schedule in the Job Schedules Table, any optional job arguments (jobargs) for the associated job are displayed in this area.

Figure 13-10 The Job Arguments Area of the Job Scheduler View



Lock	Argument	Value
<input type="checkbox"/>	agent	
<input type="checkbox"/>	image	
<input type="checkbox"/>	instance	
<input type="checkbox"/>	operation	
<input type="checkbox"/>	renameAgent	
<input type="checkbox"/>	server	

The jobargs are defined by the deployed job. Some jobs might already have a default value displayed, but others must have values specified in order for the job to be able to run.

IMPORTANT: Job arguments displayed in blue are required. You must supply data in the accompanying fields.

A job argument defines the values that can be passed in when a job is invoked. These values let you statically define and control job behavior. To learn more about each job argument, mouse over each jobarg line to display tool tip text.

The Job Scheduler uses the values you enter into the fields of this area to build a jobargs namespace in the policy for this job.

Each job argument has an accompanying *Lock* check box. When *Lock* is not selected, the accompanying job argument uses the default value specified in the job's policy. When *Lock* is selected, the value specified in the field is locked down and overrides the default value in the policy. A locked value continues to be used even if the policy value is modified.

You can click *Restore Jobargs* to restore job arguments to the values specified in the job policy. This function removes any changes you might have specified in the Job Scheduler and deselects all *Lock* check boxes.

For more information, see “[Job Arguments and Parameter Lists in Policies](#)” in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*.

User Environment

This tab displays an area (in the lower left corner of the Job Schedule Editor) that includes the *Pass User Environment* check box. Select this check box if you want to pass the assigned user's environment variables to the job when it runs. When environment variables are recorded on the user account, selecting the *Pass User Environment* check box makes those environment variables available to the job and joblet.

A user's environment is recorded under the `user.env` fact on his or her account. This fact can be set when a user logs in to PlateSpin Orchestrate and is persisted until changed. A user's environment variables can be uploaded with the `zos` command line tool at login time in one of two variations:

```
zos login --user=foo --env
```

This command uploads the entire environment to the Job Scheduler. The upload can also be seen on the User object in the Orchestrate Development Client.

```
zos login --user=foo --env=PATH
```

When the user logs in, he or she can specify one or more environment variables to use at login. The example above would result in just the PATH environment variable being uploaded.

Multiple environment variables can be specified by delimiting with a comma, as in the following example:

```
--env=PATH,LD_PATH,ID
```

NOTE: The user's environment variables can also be passed to the server when the user implements the `zos` command line tool when running a job (as opposed to logging in). The command passes the environment variable only for that particular job run.

```
zos run jobname --env=environment_variable
```

Constraints

This tab displays a constraint editor that you can use to create additional constraints for the job being scheduled. Typically, additional “resource constraints” (such as “start”) are useful to delay the start of a job when it is triggered.

Any XML constraints listed in this tab needs to have a top-level XML `<constraints>` element, as in the following example:

```
<constraints>
  <constraint type="resource">
    <contains fact="resource.groups" value="myResourceGroup" />
  </constraint>
</constraints>
```

For more information about working with constraints, see “[Scheduling with Constraints](#)” in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*.

13.1.3 Understanding Cron Syntax in the Job Scheduler

The cron triggers you can configure in the PlateSpin Orchestrate Job Scheduler use a Quartz crontrigger class for deciding when to invoke job execution. This is based on the standard Quartz format that you can find further described on the [OpenSymphony](http://www.opensymphony.com/quartz/wikidocs/CronTriggers%20Tutorial.html) (<http://www.opensymphony.com/quartz/wikidocs/CronTriggers%20Tutorial.html>) Web site, or the *KickJava* (<http://kickjava.com/src/org/quartz/CronTrigger.java.htm>) Web site.

This section includes the following information:

- ♦ “Format” on page 159
- ♦ “Special Characters” on page 159
- ♦ “Examples of Cron Syntax” on page 161
- ♦ “Cron Scheduling Precautions” on page 162

Format

A cron expression is a string comprised of 6 or 7 fields separated by white space. Fields can contain any of the allowed values, along with various combinations of the allowed special characters for that field. The fields are explained in the following table:

Table 13-2 *Fields in a Cron Expression*

Field Name	Mandatory?	Allowed Values	Allowed special Characters
Seconds	Yes	0-59	, - * /
Minutes	Yes	0-59	, - * /
Hours	Yes	0-23	, - * /
Day of the Month	Yes	1-31	, - * ? / L W
Month	Yes	1-12 or JAN-DEC	, - * /
Day of the Week	Yes	1-7 OR SUN-SAT	, - * ? / L #
Year	No	EMPTY, 1970-2099	, - * /

So cron expressions can be as simple as this:

* * * * ? *

Or cron expressions can be more complex, like this:

0 0/5 14,18,3-39,52 ? JAN,MAR,SEP MON-FRI 2002-2010

Special Characters

Cron syntax incorporates logical operators, special characters that perform operations on the values provided in the cron fields.

Table 13-3 *Special Characters in PlateSpin Orchestrate Cron Syntax*

Operator	Purpose	Example
asterisk (*)	Specifies all possible values for a field	An asterisk in the hour time field is equivalent to “every hour.”
question mark (?)	A question mark (?) is allowed in the day-of-month and day-of-week fields. It is used to specify “no specific value,” which is useful when you need to specify something in one of these two fields, but not in the other.	If you want a trigger to fire on a particular day of the month (for example, the 10th), but you don't care what day of the week that is, enter 10 in the day-of-month field, and ? in the day-of-week field. See the examples below for clarification.
dash (-)	Specifies a range of values	2-5, which is equivalent to 2, 3, 4, 5
comma (,)	Specifies a list of values	1, 3, 4, 7, 8

Operator	Purpose	Example
slash (/)	Used to skip a given number of values	<p>* / 3 in the hour time field is equivalent to 0, 3, 6, 9, 12, 15, 18, 21. The asterisk (*) specifies “every hour,” but the / 3 means only the first, fourth, seventh.</p> <p>You can use a number in front of the slash to set the initial value. For example, 2 / 3 means 2, 5, 8, 11, and so on.</p>
L (“last”)	<p>The L character is allowed for the day-of-month and day-of-week fields.</p> <p>Specifies either the last day of the month, or the last xxx day of the month.</p>	<p>The value L in the day-of-month field means “the last day of the month”—day 31 for January, day 28 for February in non-leap years. If you use L in the day-of-week field by itself, it simply means 7 or SAT. But if you use it in the day-of-week field after another value, it means “the last xxx day of the month.” For example, 6L means “the last Friday of the month.”</p> <hr/> <p>TIP: When you use the L option, be careful not to specify lists or ranges of values. Doing so causes confusing results.</p>
W (“weekday”)	<p>The W character is allowed for the day-of-month field.</p> <p>Specifies the weekday (Monday-Friday) nearest the given day.</p>	<p>If you specify 15W as the value for the day-of-month field, the meaning is “the nearest weekday to the 15th of the month.” So if the 15th is a Saturday, the trigger fires on Friday the 14th. If the 15th is a Sunday, the trigger fires on Monday the 16th. If the 15th is a Tuesday, it fires on Tuesday the 15th. However, if you specify 1W as the value for day-of-month, and the 1st is a Saturday, the trigger fires on Monday the 3rd, because it does not “jump” over the boundary of a month’s days. The W character can only be specified when the day-of-month is a single day, not a range or list of days.</p> <hr/> <p>TIP: You can combine the L and W characters for the day-of-month expression to yield LW, which translates to “last weekday of the month.”</p>
pound sign (#)	The pound sign (#) character is allowed for the day-of-week field. This character is used to specify “the nth” xxx day of the month.	<p>The value of 6#3 in the day-of-week field means the third Friday of the month (day 6 = Friday and #3 = the 3rd one in the month).</p> <p>Other Examples: 2#1 specifies the first Monday of the month and 4#5 specifies the fifth Wednesday of the month. However, if you specify #5 and there are fewer than 5 of the given day-of-week in the month, no firing occurs that month.</p>

NOTE: The legal characters and the names of months and days of the week are not case sensitive. MON is the same as mon.

You can specify days in two fields: month day and weekday. If both are specified in an entry, they are cumulative, meaning that both of the entries are executed.

Examples of Cron Syntax

The following table shows examples of full cron expressions and their respective meanings.

Table 13-4 Results of Altered Cron Syntax on Execution Times

Cron Expression Example	Description
0 0 12 * * ?	Fire at 12:00 p.m. (noon) every day
0 15 10 ? * *	Fire at 10:15 a.m. every day
0 15 10 * * ?	Fire at 10:15 a.m. every day
0 15 10 * * ? *	Fire at 10:15 a.m. every day
0 15 10 * * ? 2008	Fire at 10:15 a.m. every day during the year 2008
0 * 14 * * ?	Fire every minute starting at 2:00 p.m. and ending at 2:59 p.m., every day
0 0/5 14 * * ?	Fire every five minutes starting at 2:00 p.m. and ending at 2:55 p.m., every day
0 0/5 14,18 * * ?	Fire every five minutes starting at 2:00 p.m. and ending at 2:55 p.m., and fire every five minutes starting at 6:00 p.m. and ending at 6:55 p.m., every day
0 0-5 14 * * ?	Fire every minute starting at 2:00 p.m. and ending at 2:05 p.m., every day
0 10,44 14 ? 3 WED	Fire at 2:10 p.m. and at 2:44 p.m. every Wednesday in the month of March
0 15 10 ? * MON-FRI	Fire at 10:15 a.m. every Monday, Tuesday, Wednesday, Thursday and Friday
0 15 10 15 * ?	Fire at 10:15 a.m. on the 15th day of every month
0 15 10 15 * ?	Fire at 10:15 a.m. on the last day of every month
0 15 10 ? * 6L	Fire at 10:15 a.m. on the last Friday of every month
0 15 10 ? * 6L 2008-2011	Fire at 10:15 a.m. on every last Friday of every month during the years 2008, 2009, 2010, and 2011
0 15 10 ? * 6#3	Fire at 10:15 a.m. on the third Friday of every month
0 0 12 1/5 * ?	Fire at 12:00 p.m. (noon) every five days every month, starting on the first day of the month
0 11 11 11 11 ?	Fire every November 11th at 11:11 a.m.

Cron Scheduling Precautions

You should remember the following items when you use cron scheduling:

- ♦ Always check the effect of adding the ? and * characters in the day-of-week and day-of-month fields to make sure the expected behavior fires correctly.
- ♦ Support for specifying both a day-of-week and a day-of-month value is not complete (you must currently use the ? character in one of these fields).
- ♦ Be careful when setting fire times to occur between 12:00 a.m. and 1:00 a.m.— changing to or out of Daylight Saving Time can cause a skip or a repeat in the schedule firing, depending on whether the clock moves backward or forward.

13.2 Walkthrough: Scheduling a System Job

This section demonstrates how you can use the Orchestrate Development Client to deploy and schedule an existing system job named `auditcleaner.job`. This example job is included in the `examples` directory of your PlateSpin Orchestrate installation.

This section includes the following information:

- ♦ [Section 13.2.1, “Deploying a Sample System Job,” on page 162](#)
- ♦ [Section 13.2.2, “Creating a New Schedule for the Job,” on page 164](#)
- ♦ [Section 13.2.3, “Defining the New Schedule,” on page 165](#)
- ♦ [Section 13.2.4, “Activating the New Schedule,” on page 172](#)
- ♦ [Section 13.2.5, “Running the New Schedule Immediately,” on page 172](#)

13.2.1 Deploying a Sample System Job

Before a job can run, the PlateSpin Orchestrate administrator must deploy it, which involves moving it from a developed package state to a state where it is ready and available for users. Only the administrator has the necessary rights to deploy a job.

There are four methods you can use to deploy a job:

- ♦ Deploy it from the PlateSpin Orchestrate Development Client by right-clicking the *Jobs* container in the Explorer panel.
- ♦ Deploy it from the PlateSpin Orchestrate Development Client by selecting the *Actions* menu in the Orchestrate Development Client.
- ♦ Deploy it from the `zosadmin` command line (`zosadmin deploy path_to_job`).
- ♦ Copy the deployable component to the “hot” deployment directory under the Orchestrate Server instance directory. Typically, this directory is located at `/var/opt/novell/zenworks/zos/server/deploy`. Using this method, deployment proceeds within a few seconds. PlateSpin Orchestrate monitors this directory.

A runnable job can also be scheduled, which means that the schedule for running the job and the trigger or triggers that initiate or “fire” the schedule (or both) are configured and packaged with the job.

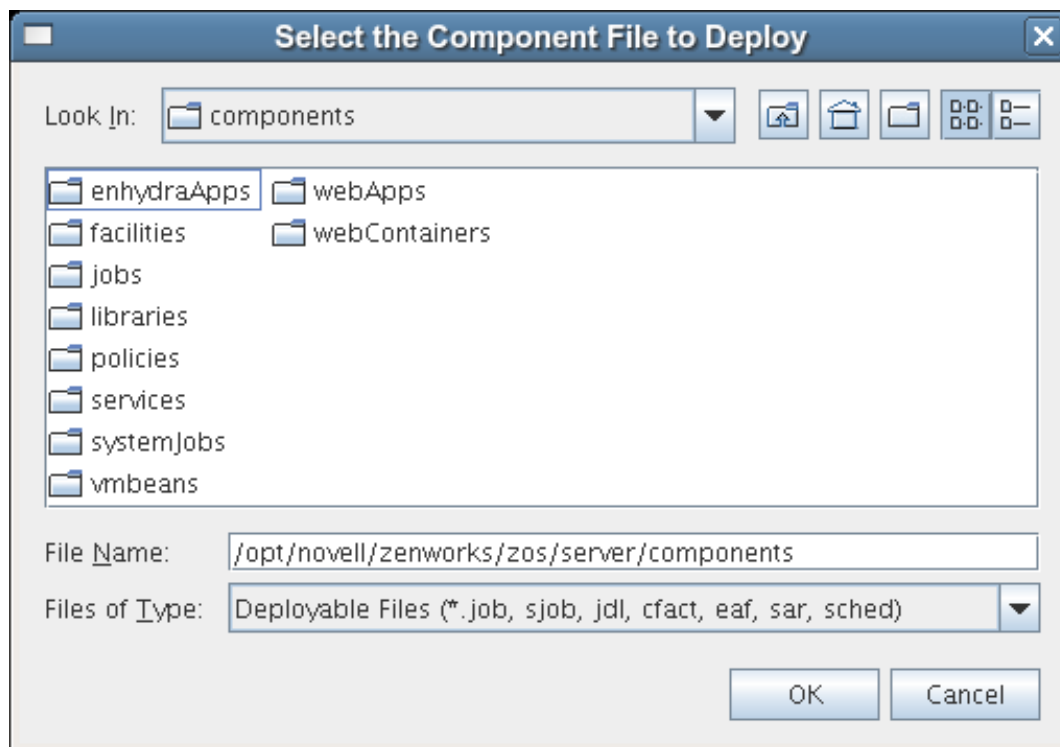
For this walkthrough, you deploy one of several system jobs (`auditCleaner.job`) developed for PlateSpin Orchestrate customers to demonstrate how system jobs are deployed and run. This job package, which is actually a `.jar` archive, includes only a `.policy` component and a `.jdl` component. It does not have a `.sched` component. You can use the Job Scheduler in the Orchestrate Development Client to add the `.sched` component separately.

NOTE: A PlateSpin Orchestrate job developer can create and package jobs that include a `.jdl` file, a `.policy` file, a `.trig` file (trigger), and a `.sched` file (schedule). The presence of the `.sched` file in the job package is also typical of the predeployed discovery jobs installed with PlateSpin Orchestrate, which run without intervention when the criteria for firing the schedule are satisfied. Such jobs are visible in the Job Scheduler because they already include the `.sched` components.

For more information about developing jobs and schedules in a job archive, see “[Job Scheduling](#)” in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*.

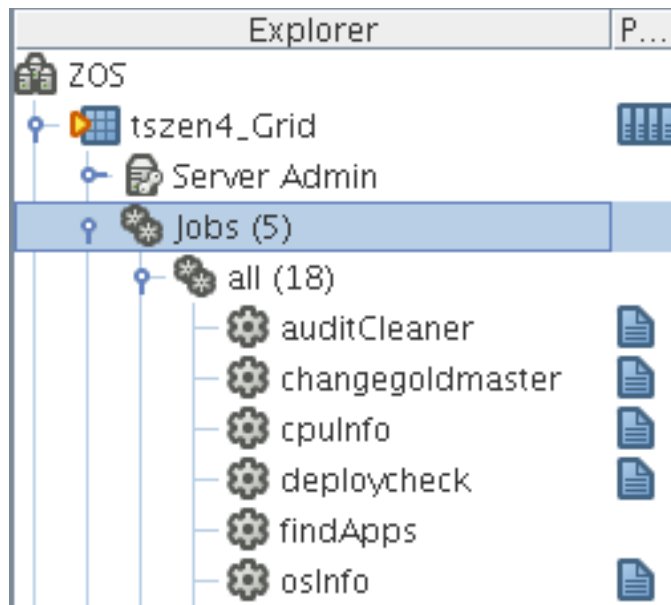
Although this walkthrough demonstrates only the first method listed above for deploying, the other methods are relatively simple, so no further examples are provided to illustrate them.

- 1 In the Explorer panel of the PlateSpin Orchestrate Development Client, right-click the *Jobs* container, then click *Deploy Job* to open the Select the Component File to Deploy dialog box.




- 2 Open the *Look In* drop-down list, then navigate to the location of the job you want to deploy. Although a job developer can store PlateSpin Orchestrate jobs at any location on the network, the sample jobs shipped with PlateSpin Orchestrate are limited to the directories where the product is installed. For this walkthrough, navigate to the `/opt/novell/zenworks/zos/server/components/systemJobs` directory.
- 3 Select *auditCleaner.job*, then click *OK* to deploy the job to the *Jobs* container.

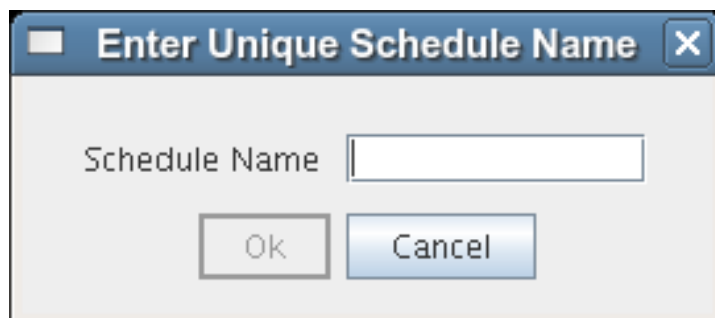
The job appears in the *all* container and in the *examples* container in the tree.





13.2.2 Creating a New Schedule for the Job

When a job has been deployed, you can create a schedule to specify when you want it to run. In this walkthrough, you create a schedule for the `auditCleaner` job by using the Scheduler tool in the Orchestrator Development Client.

- 1 From the toolbar in the Orchestrator Development Client, click the Job Scheduler icon  to open the Job Scheduler view.
- 2 In the Job Scheduler view, click *New* to open the Enter Unique Schedule Name dialog box.



- 3 Specify a name for the schedule you want to create for this job. For this walkthrough, specify the name `cleaner` in the *Schedule Name* field, then click *OK* to return to the Job Scheduler view.

 Schedule Name	Job Name	Priority	User ID	Status	Last Job ID	Last Job Status	Last Fire Time	Active Jobs
 cleaner				 Paused		Not fired yet	n/a	

The new schedule is highlighted in the Job Schedules Table and is flagged with a pencil icon, signifying that the schedule has not been committed yet. Continue with [Section 13.2.3, “Defining the New Schedule,” on page 165](#) to define this new schedule by adding the specific information you want.

13.2.3 Defining the New Schedule


Defining a new job schedule consists of selecting its general properties, its specific properties, and the triggers you want to be associated with it. This section includes the following information:

- ♦ [“Choosing General Properties for a New Schedule” on page 165](#)
- ♦ [“Creating and Assigning a Time Trigger for the New Schedule” on page 167](#)
- ♦ [“\(Optional\) Adding Specific Parameters to the New Schedule” on page 170](#)


Choosing General Properties for a New Schedule

After you have created a new job schedule, its name cannot be changed, but you can add properties to it that help to identify and classify it in a general way. Use the following steps to add these properties:

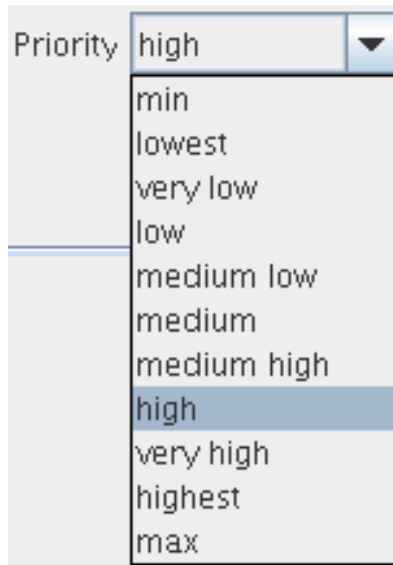
- 1 In the Job Schedule Editor panel of the Scheduler view, select the *Job* drop-down list.



- 2 From the list of available jobs, select *auditCleaner* as the job to which this schedule applies.
- 3 In the Job Schedule Editor, select the *User* drop-down list.




- 4 From the list of available users, select *zosSystem* as the user who runs this job.
The *zosSystem* user is the built-in user that is always present. It is commonly used for routine jobs like this example.
- 5 In the Job Schedule Editor, select the *Priority* drop-down list.



- 6 From the list of available priorities, select *high* as the priority for this job schedule.

The maximum selectable priority is dependent on an attribute associated with the selected user.

- 7 Click the *Save* icon  on the toolbar of the Orchestrate Development Client to save the general properties you have selected for the new schedule.

The schedule is now committed, and the attribute columns in the Job Schedules Table are populated with the name of the job that the schedule will run, the user it will run as, the priority at which it will run, and its current status. Because the schedule has not been activated yet, it remains in a *Disabled* state.

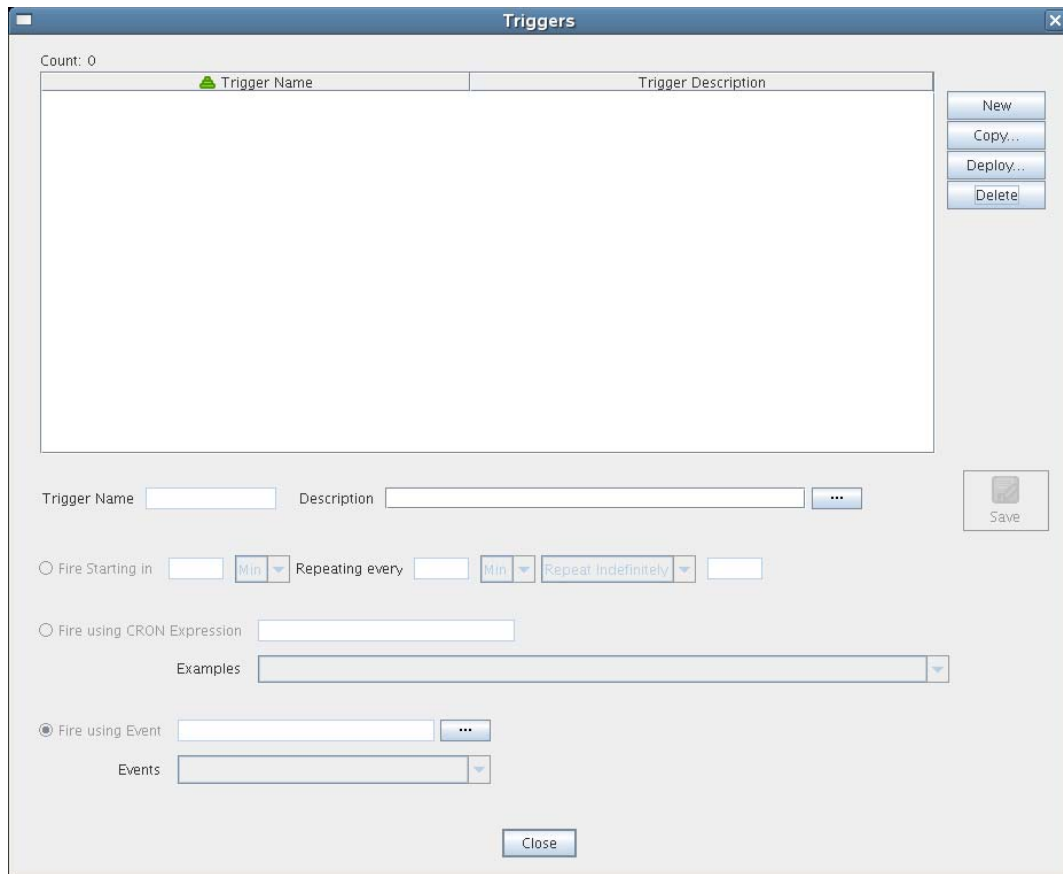
When you have chosen the general properties of the new schedule, you can either continue by [“\(Optional\) Adding Specific Parameters to the New Schedule” on page 170](#) or by proceeding directly to [“Creating and Assigning a Time Trigger for the New Schedule” on page 167](#).

Creating and Assigning a Time Trigger for the New Schedule

A job already defined in a schedule can be triggered with two main themes: the occurrence of an event or the arrival of a point in time. In this walkthrough, you define a time trigger for the `cleaner` schedule.

In this example, there are no defined time triggers in the Job Scheduler, so use the following steps to define a time trigger.


- 1 In the Job Schedule view, click *Edit Triggers* to display the Triggers dialog box.



Time triggers are shareable across schedules. After a time trigger is defined, it is added to a list of triggers in this dialog box. You can select a predefined trigger from this list when you create a new schedule, or you can create a new time trigger, as the next steps demonstrate.

NOTE: For detailed information about cron syntax, see [“Understanding Cron Syntax in the Job Scheduler” on page 158](#).

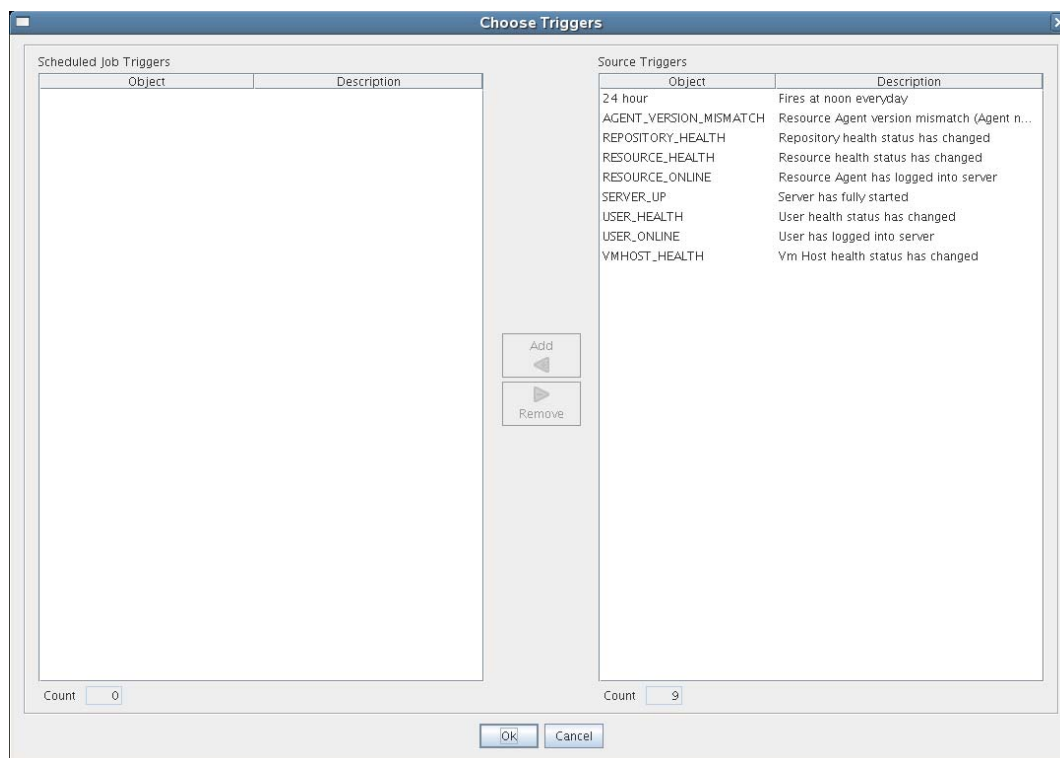
- 2 In the Triggers dialog box, click *New* to clear and activate the fields in the dialog box for the creation of a unique time trigger.
- 3 In the Enter Unique Trigger Name dialog box, specify `24 hour` as the unique name of this time trigger, then click *OK*.
- 4 In the *Description* field, specify `Runs every 24 hours at noon` as the description for this time trigger.
- 5 Click *Fire Using CRON Expression* to activate the fields for defining a cron expression.



- 6 Click the drop-down list of sample cron expressions, then select the default cron expression, `0 0 12 * * ?`, which is listed first.

The sample expressions in the drop-down list show cron strings with accompanying descriptions to remind you how a cron string is constructed. The examples are selectable and editable and can be used in the schedule, just as you have done in this step.

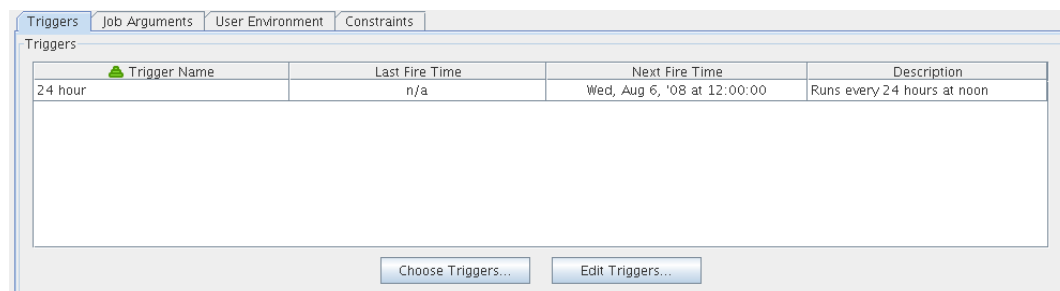
- 7 Click *Save* to save the trigger you just created, then click *Close* to return to the Job Scheduler view.
- 8 From the Job Scheduler view, make sure that the *cleaner* schedule is selected, then click *Choose Triggers* to display the Choose Triggers dialog box.




- 9 In the Choose Triggers dialog box, select *24 hour* (the name of the trigger you created), click *Add* to move the trigger definition to the *Scheduled Job Triggers* list, then click *OK* to return to the Job Scheduler view.

NOTE: You can select and combine as many time triggers as you want to apply to a given schedule. You can also combine time triggers with event triggers on a given schedule.

In the *Triggers* list of the Job Scheduler view, the *24 hour* trigger is now associated with the new schedule.



- 10 Click the *Save* icon  to update the Orchestrate Server with the new schedule/trigger association.

(Optional) Adding Specific Parameters to the New Schedule

You can now add specific parameters to the schedule to edit its job arguments, to choose whether you want to pass the user environment variables to the schedule, or to specify policy constraints to further focus the purpose of this schedule when it fires.

For the purpose of this walkthrough, none of these specific parameters is modified, although a general overview of how to do so is explained.

The following specific parameters can be managed in the Job Scheduler Editor:

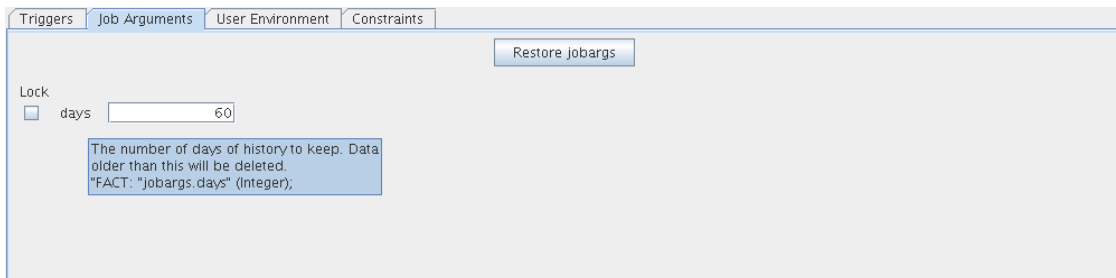
- ♦ “Job Arguments” on page 170
- ♦ “User Environment” on page 170
- ♦ “Constraints” on page 171

Job Arguments

As explained in [Section 13.1.2, “Creating or Modifying a Job Schedule,” on page 148](#), a job argument defines the values that can be passed in to the process when a job is invoked. These values let you statically define and control job behavior. The job arguments that appear in the *Job Arguments* tab of the Schedule Editor depend on the job. The job might have no arguments.

By default, the auditCleaner job lists only one job argument, *jobargs.days*.

Figure 13-11 The Job Arguments Tab of the Job Schedule Editor



According to the tool tip text, this argument is the number of days of job history to keep, so this job cleans up the history of the job in the PlateSpin Orchestrate audit database after the job reaches the age of 60 days. Data older than 60 days is to be deleted. If you want to, you can change this parameter, or any other parameter in a job argument.

If the default value for a job argument parameter is missing, the job might fail, so you should inspect these parameters carefully.

User Environment

As explained in [Section 13.1.2, “Creating or Modifying a Job Schedule,” on page 148](#), a user’s environment variables are available in the Job Scheduler only if that user utilizes the *zos* command line tool and elects to pass those environment variables to the server at login time or when he or she runs a job (running the job creates the environment variables as facts in the job). The *zos run* command passes the environment for that particular job run only.

In this walkthrough, the *zosSystem* user shows no user environment variables.

Figure 13-12 The User Environment Tab of the Job Scheduler Editor, No User Environment Variables Available

The screenshot shows the 'User Environment' tab of the Job Scheduler Editor. At the top, there are four tabs: 'Triggers', 'Job Arguments', 'User Environment', and 'Constraints'. The 'User Environment' tab is selected. Below the tabs, there is a checkbox labeled 'Pass User Environment' which is unchecked. Below the checkbox, the text 'No User Environment Variables' is displayed.

Because there are no environment variables listed, there are none to pass to the schedule, so it is not necessary to select the *Pass User Environment* check box. By default, this check box is not selected, even if environment variables are present for a user specified to run the job.

Sometimes a job is written to work from a user's environment variables. In this case, if a user has not logged in or has not run the job from the zos command line using the necessary environment option, the schedule must pass those variables to the job when it is invoked.

If you associated a user who had user environment variables with this schedule, you would see a list of those environment variables as they would be passed to the job.

Figure 13-13 The User Environment Tab of the Job Schedule Editor, User Environment Variables Available

The screenshot shows the 'User Environment' tab of the Job Scheduler Editor. At the top, there are four tabs: 'Triggers', 'Job Arguments', 'User Environment', and 'Constraints'. The 'User Environment' tab is selected. Below the tabs, there is a checkbox labeled 'Pass User Environment' which is unchecked. Below the checkbox, a list of environment variables is displayed in a table-like format:

FROM_HEADER	
XFILESEARCHPATH	/usr/dt/app-defaults/%L/Dt
LANG	en_US.UTF-8
G_FILENAME_ENCODING	@locale,UTF-8,ISO-8859-15,CP1252
-	/opt/novell/zenworks/zos/clients/.../jre/bin/java
MORE	-sl

Selecting the *Pass User Environment* check box in this scenario would create these variables as facts used for this job invocation.

Constraints


As explained in [Section 13.1.2, "Creating or Modifying a Job Schedule," on page 148](#), the *Constraints* tab displays a constraint editor that you can use to create additional constraints for the job being scheduled.

Figure 13-14 The Constraints Tab of the Job Schedule Editor

The screenshot shows the 'Constraints' tab of the Job Scheduler Editor. At the top, there are four tabs: 'Triggers', 'Job Arguments', 'User Environment', and 'Constraints'. The 'Constraints' tab is selected. Below the tabs, there is a list of constraints. The list is currently empty, showing only a vertical column of numbers from 1 to 11, indicating the positions for constraints.

Any other constraints associated with the context of this job invocation (including but not limited to this job), with the user you've selected, with that user's group, with the jobs group, with the resources that the job uses, or with the resource groups that the job uses, run in spite of the policy that you define here. These additional constraints usually restrict or refine what the job does when this schedule fires.

These constraints are passed to the job only when this schedule is invoked. For example, you could add a start constraint to delay the start of a job, a resource constraint to run on only one of three named machines, or a continue constraint to automatically time out the job if it takes too long to run. Anything you can do with a regular job policy constraint, you can add as a special constraint here for this particular schedule invocation.







Click *Save* icon  to update the Orchestrate Server with the new schedule.

For more information about policies, see “[Policies](#)” in the [PlateSpin Orchestrate 2.5 Developer Guide and Reference](#).

13.2.4 Activating the New Schedule

When the new schedule has been created and its triggers defined, you need to take it from the disabled state to an active state where it is ready to run.

- 1 In the Job Scheduler view, select the newly created job. The job shows that it is in a *Disabled* state.

 Schedule Name	 Job Name	Priority	User ID	Status	Last Job ID	Next Fire Time	Last Fire Time
 cleaner	 auditCleaner	high	 zosSystem	 Disabled		12:00:00	n/a

- 2 Click *Enable* to enable the schedule.



 Schedule Name	 Job Name	Priority	User ID	Status	Last Job ID	Next Fire Time	Last Fire Time
 cleaner	 auditCleaner	high	 zosSystem	 Enabled		12:00:00	n/a




The schedule is now enabled, but has not run yet.



13.2.5 Running the New Schedule Immediately

You can trigger the schedule immediately, rather than waiting for the triggers to fire.


- 1 In the Job Schedules Table of the Job Scheduler view, select *cleaner* (the name of the schedule you want to run), click *Run Now*, then click the job monitor icon on the toolbar (*Jobs*) to open the Job Monitor view.






 Submit Time	Job ID	Instance Name
16:52:13	 zosSyst...	Scheduler(cleaner)

 Joblets	 Resources	 Policy Debugger
---	---	---


 Status: 1 Joblet
Memo:








The joblet icon shows that the job is running.

- 2 Click the Job Scheduler icon  on the toolbar to open the Job Scheduler view.

 Schedule Name	Job Name	Priority	User ID	Status	Last Job ID	Last Fire Time	Active Jobs
 cleaner	 auditCleaner	high	 zosSyst...	 Enabled		12:44:49	zosSystem.audit...

The `cleaner` schedule is listed as an active job. This indicates that the schedule has started the job as anticipated.

If you click the refresh icon , you can see that the job now has a Job ID.

 Schedule Name	Job Name	Priority	User ID	Status	Last Job ID	Last Fire Time	Last Job Status
 cleaner	 auditCleaner	high	 zosSystem	 Enabled	zosSystem.auditCleaner.9	12:44:49	 Job failed because of...

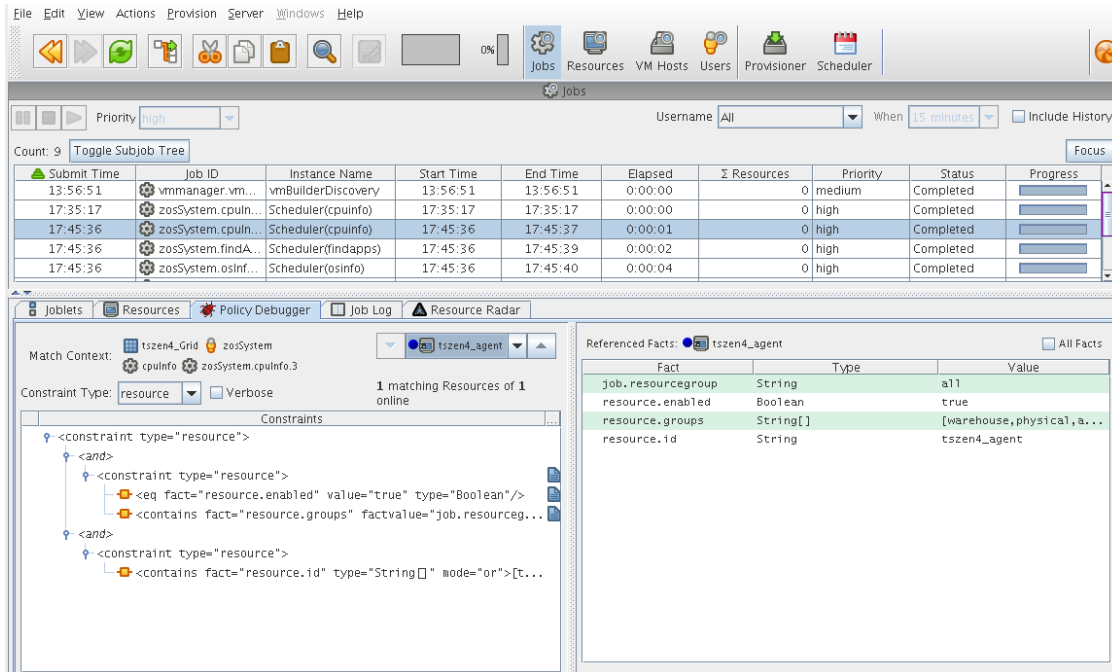
If the job invocation fails, as in this example, a red exclamation icon is also displayed.

The Policy Debugger

14

The Policy Debugger is a tabbed page available in several views of the PlateSpin Orchestrate Development Client. This tool helps you to determine the reasons for the current state of a job. The following figure shows the Policy Debugger tab opened in the Jobs view of the Orchestrate Development Client.

Figure 14-1 PlateSpin Orchestrate Jobs View with the Policy Debugger Page Open



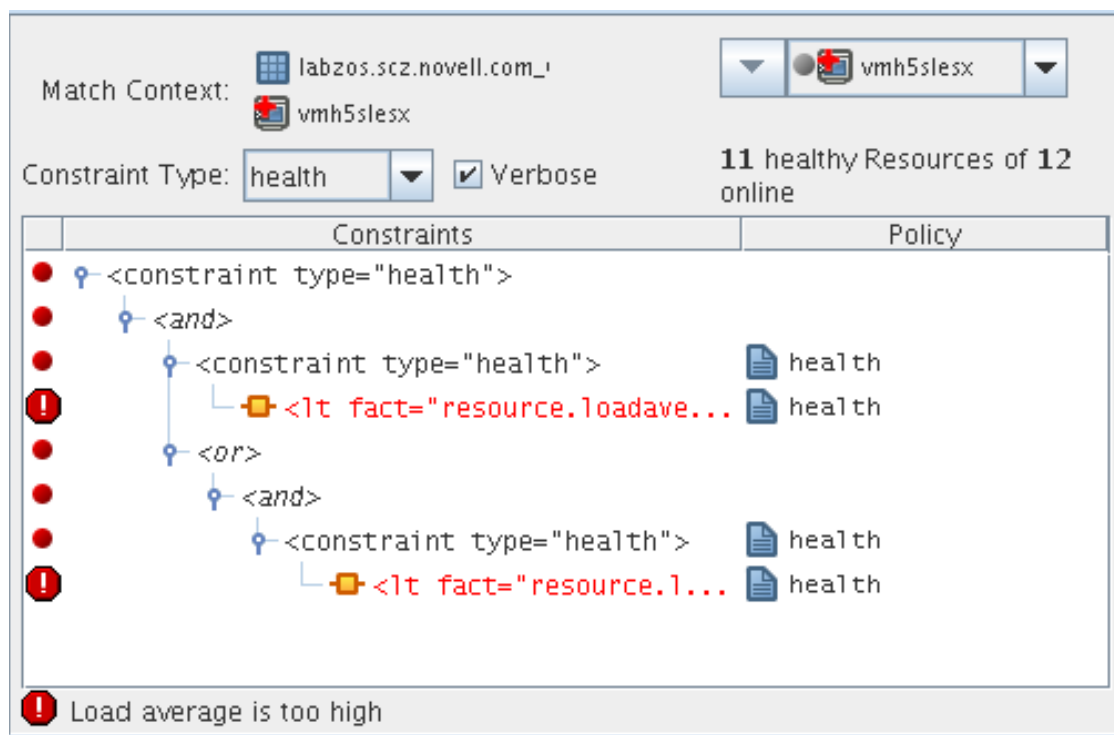
The Policy Debugger tab is also available in the VM Hosts view and in the Provisioner view of the Orchestrate Development Client.

- ♦ [Section 14.1, “The Constraints Table View,” on page 175](#)
- ♦ [Section 14.2, “The Facts Table View,” on page 180](#)
- ♦ [Section 14.3, “Policy Debugger Use Cases,” on page 182](#)

14.1 The Constraints Table View

The left side of the Policy Debugger page is referred to as the Constraints Table view.

Figure 14-2 The Constraints Table View



The appearance of this view can change, depending on the constraint type you select in the drop down menu. For a description of these types, see [Section 14.1.2, “The Constraint Type List,” on page 178](#).


The Constraints Table View is composed of several parts:

- ♦ [Section 14.1.1, “The Match Context Area,” on page 176](#)
- ♦ [Section 14.1.2, “The Constraint Type List,” on page 178](#)
- ♦ [Section 14.1.3, “The Verbose Check Box,” on page 178](#)
- ♦ [Section 14.1.4, “The Capable Resources Summary,” on page 178](#)
- ♦ [Section 14.1.5, “The Constraints Column of the Constraints Table View,” on page 178](#)
- ♦ [Section 14.1.6, “The Policy Column of the Constraints Table,” on page 180](#)





14.1.1 The Match Context Area

The policy debugger provides the general identification of a job instance in the *Match Context* area of the Constraints Table View. The Match Context defines everything associated with running a job on a particular resource it references Facts, which are also referenced in Policies. The Policies define how, when, and where the job runs.

Figure 14-3 The Match Context Area of the Constraints Table View in the Policy Debugger

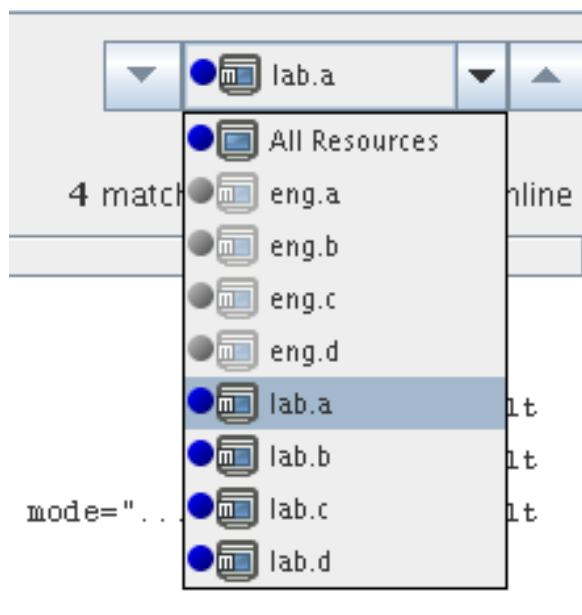


That identifying Facts in the Match Context include:

- ♦ Matrix: The  icon and a text string identifies the machine that matches the grid name given to the Orchestrate Server where this job is running.
- ♦ User: The  icon and a text string identifies the User object that matches the user who is running the job.
- ♦ Job: The  icon and a text string identifies the deployed Job that matches the one that is running on the grid.
- ♦ Job Instance: The  icon and a fully qualified text string identifies the specific Job instance that matches the deployed job running on the grid.
- ♦ Resource: The Resource drop down list shows all resources. The list appears in the Match Context if the *resource* constraint type is selected. The resources in the list that are currently offline display with a dimmed icon. If available, a listed resource has a colored dot by its side. The color of the dot (blue ● or gray ●) and the resource type it accompanies has significance:
 - ♦ A blue dot with the *All Resources* label indicates that at least one resource matches the constraints and is capable of servicing the job.
 - ♦ A gray dot with the *All Resources* label indicates that no resources match the constraints.
 - ♦ A blue dot with a named, selected resource indicates that its constraints match and it is capable of servicing the job.
 - ♦ A gray dot by a named, selected resource indicates that its constraints do not match and that it is not capable of servicing the job.

The following figure shows a list of eight resources. Four of those resources (*lab.a*, *lab.b*, *lab.c* and *lab.d*) are online, and their constraints match so they are capable of servicing the job.

Figure 14-4 Resource Drop Down List Showing Online and Offline Resources



14.1.2 The Constraint Type List

Select one of the constraint types in the drop down list to specify a policy context. Constraint types in the list are disabled (dimmed) if they do not apply to the job that you are debugging.

- ♦ **accept:** accept
- ♦ **start:** start
- ♦ **continue:** continue
- ♦ **provision:** provision
- ♦ **allocation:** allocation
- ♦ **resource:** resource
- ♦ **vmhost:** vmhost
- ♦ **repository:** repository
- ♦ **health:** health

14.1.3 The Verbose Check Box

When you select the Verbose check box, the `reason` string specified in the policy is displayed in the *Constraints* tree. For more information, see [Section 14.1.5, “The Constraints Column of the Constraints Table View,”](#) on page 178.

14.1.4 The Capable Resources Summary

Directly under the Resource List in the constraint view, a populated string summarizes the resources that are capable of servicing the job. For example, 4 matching Resource of 10 online indicates that four of the ten total online resources are capable of servicing the job.

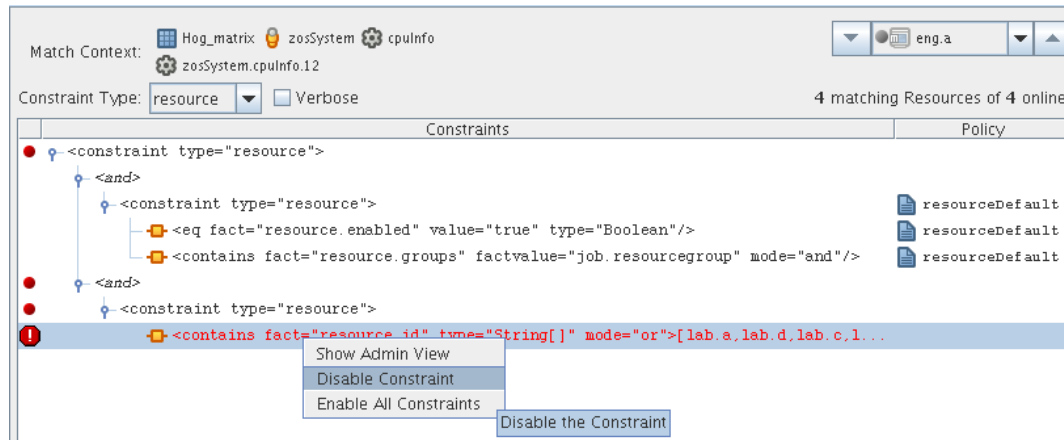
14.1.5 The Constraints Column of the Constraints Table View

The Constraints column of the constraints table view shows the logical hierarchy (that is, a “tree” format) of the constraints that are defined by the Policies associated with the Job. You can identify the status of the listed constraints by the icons that might be displayed in the far left column of the table:

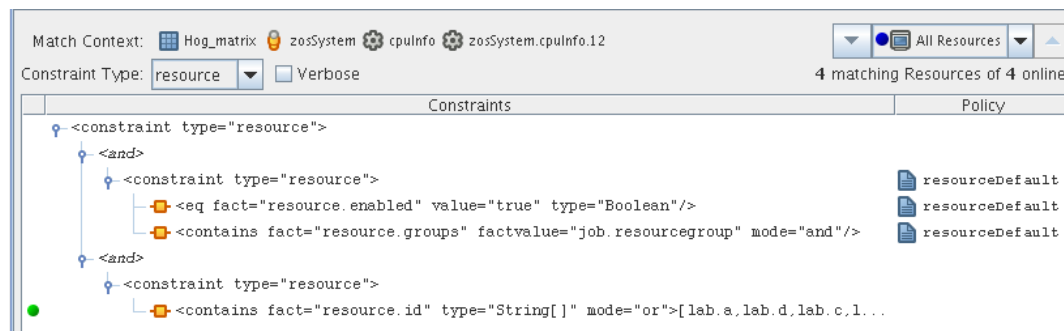
- ♦ no icon: The constraint passes the match. It is a “true” match. The figure below shows that the resource `lab.a` is available to run the job because all of its constraints match. No red icons are displayed next to any listed constraint.

Constraints	Policy
<pre><constraint type="resource"> <and> <constraint type="resource"> <eq fact="resource.enabled" value="true" type="Boolean"/> <contains fact="resource.groups" factvalue="job.resourcegroup" mode="and"/> </and> <constraint type="resource"> <contains fact="resource.id" type="String[]" mode="or">[lab.a,lab.d,lab...</pre>	resourceDefault

- ♦ red dot icon: The constraint does not pass the match. The figure below shows that the resource eng.a cannot run the job because its constraints do not match.



- ♦ red octagonal icon: The constraint does not pass the match and is blocking the job. The figure above also shows the blocking constraint (red octagon).
- ♦ green dot icon: A blocking constraint has been disabled so that it behaves like a match. The figure below shows the green dot icon next to that the constraint that was formerly blocked and can now behave as a match.



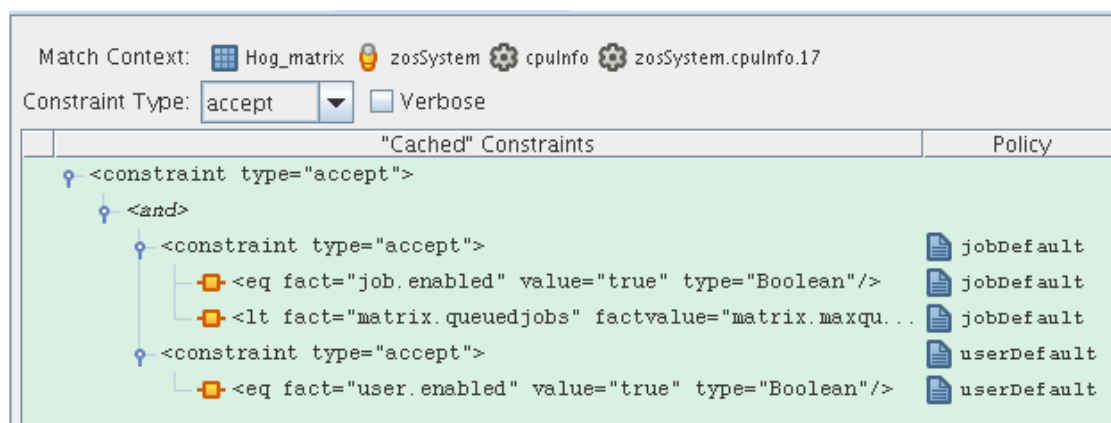
If you right-click a constraint in the table, a popup menu with three options is displayed. This menu lets you change the status of the constraint.

- ♦ **Show Admin View:** Select this option to open the Admin View for the specific resource selected.
- ♦ **Disable Constraint:** Select this option to disable (attach a green dot icon to) a constraint. Disabling a constraint with this function effectively makes it match, a condition that can prove useful if you want to perform a “what if” test without actually changing a policy.
- ♦ **Enable All Constraints:** Select this option if you have disabled one or more constraints during testing and you want to restore them to the enabled state.

Cached Constraints in the Constraints Column

When you change the constraint type in the Constraints Type List, the background of the table changes to green for some types. These are “cached” constraints that are saved with the job after it has completed. Their purpose is to help you debug the policy.

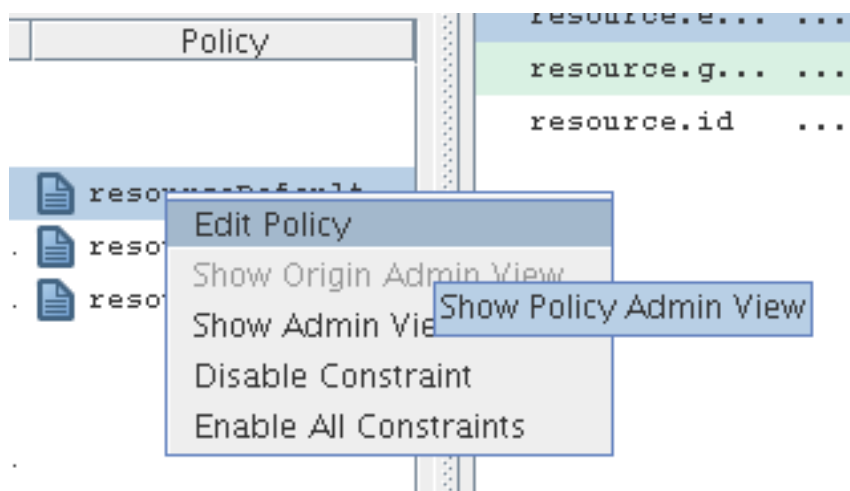
Figure 14-5 *Cached Constraints Displayed in the Constraints Table View*



14.1.6 The Policy Column of the Constraints Table

The Policy column of the constraints table displays the policy name that contributed the constraint. Right-click a policy name to open a popup menu offering the option to open the policy editor for the specified policy. The menu also includes constraint enabling or disabling options, just as the popup menu for constraint column does.

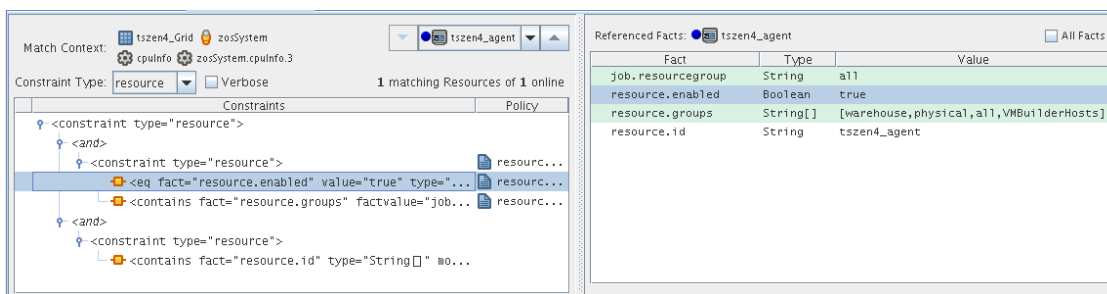
Figure 14-6 *The Popup Menu Launched from the Policy Column*



14.2 The Facts Table View

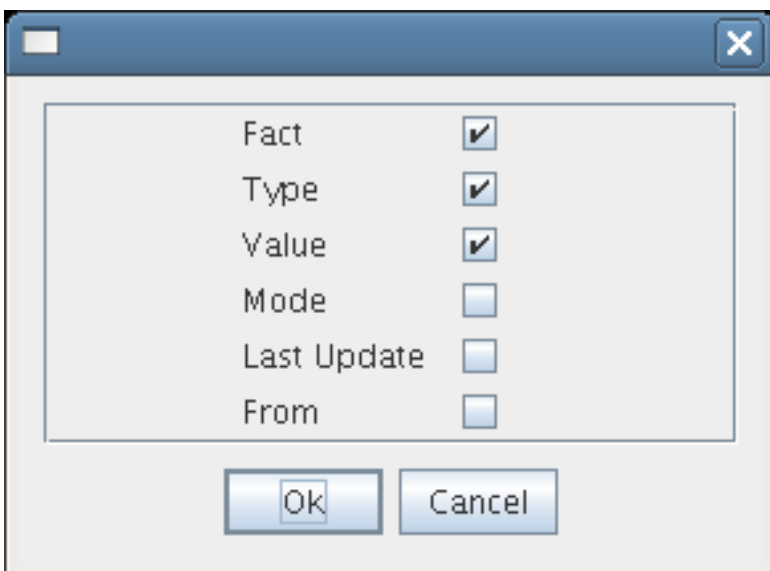
The Facts Table view displays the facts referenced in the Constraint Tree view for a specified Resource. Selecting a fact in the Constraint tree automatically selects that fact in the table.

Figure 14-7 The Constraints Table View and the Accompanying Facts Table View



If you right-click a column head in this table, a menu is launched where you can select the columns that you want to display.

Figure 14-8 Menu Used to Select the Columns Displayed in the Facts Table View of the Policy Debugger



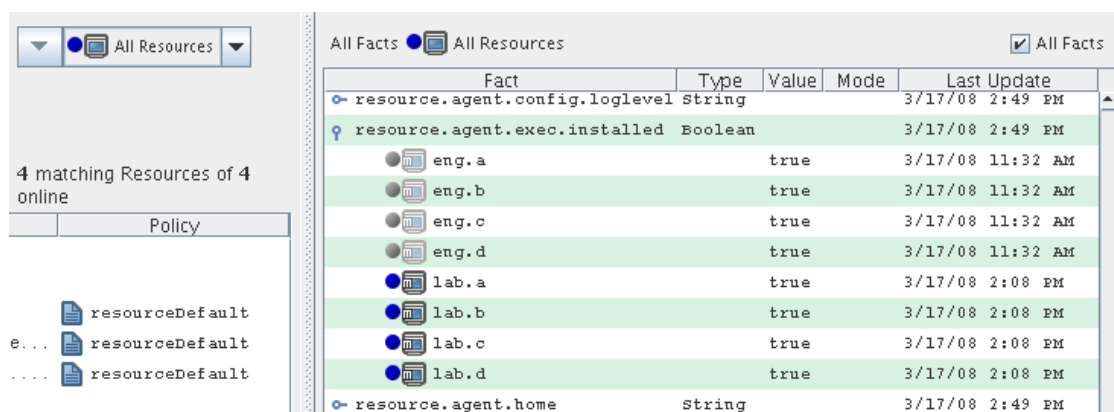
- ♦ [Section 14.2.1, “The All Facts Check Box,” on page 181](#)

14.2.1 The All Facts Check Box

If you select the *All Facts* check box at the top of the Facts Table view, all of the facts (including matrix, user, job, jobargs, jobinstance, and resource facts) associated with the Match Context are listed.

If you select *All Resources* in the Match Context (see [Section 14.1.1, “The Match Context Area,” on page 176](#)) and you also select the *All Facts* check box, the Facts Table view displays all the facts for all resources for the specified Match Context.

Figure 14-9 All Facts Check Box Selected with All Resources in Match Context



14.3 Policy Debugger Use Cases

This section includes several use cases that show how the Policy Debugger works and how to use it. The following use cases are included:

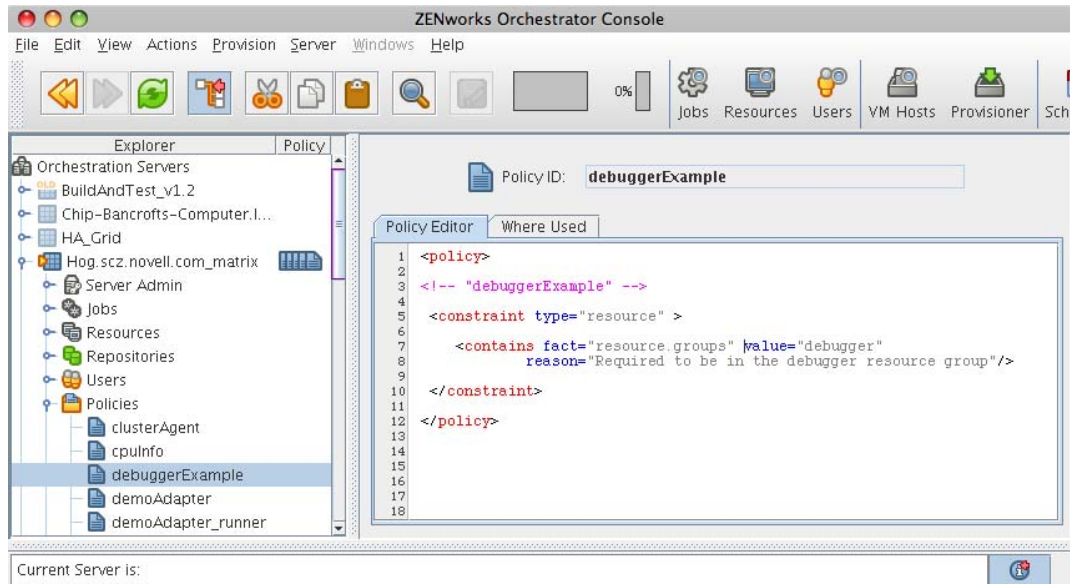
- ♦ [Section 14.3.1, “Use Case 1: Determining Why a Job is in a Waiting State,” on page 182](#)

14.3.1 Use Case 1: Determining Why a Job is in a Waiting State

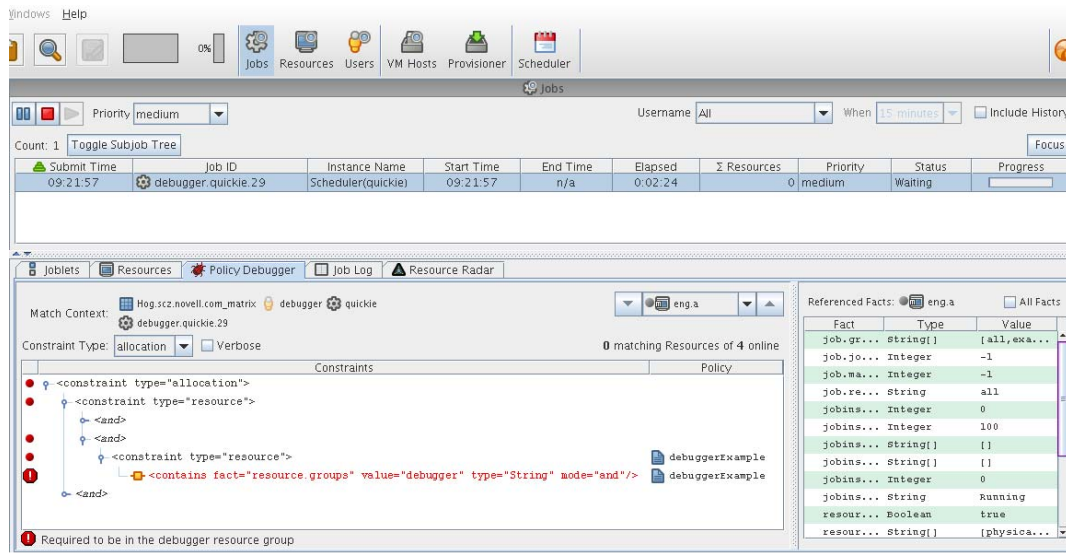
The objective of this use case is to run a job that sits in the waiting state and to use the policy debugger to identify why it is in the state and to make the necessary changes to get the job to run. The `quickie.job` is used along with a simple policy that specifies that the Resources that are to be used must be in a Resource group called `debugger`.

Use the following steps to recreate the use case.

- 1 In the Orchestrate Development Client, create a user named `debugger`.
- 2 Deploy `quickie.job` from the `/examples` directory.
- 3 In the Orchestrate Development Client, create a schedule named `quickie`, specifying the `quickie` job and the `debugger` user.
- 4 In the Orchestrate Development Client, create a policy and name it `debuggerExample`. The policy needs to specify that the resource used belongs to the group called `debugger`.



- 5 In the Orchestrate Development Client, associate the `debuggerExample` policy to the `quickie` job.
- 6 In the Job Scheduler view of the Orchestrate Development Client, select the `quickie` schedule, then click *Run Now* to run the `quickie` schedule.
- 7 In the Job Monitor view of the Orchestrate Development Client, select the *Policy Debugger* tab and verify that the job is in the waiting state.
- 8 In the Constraints Table view, open the *Constraint Type* drop down list, then select *Allocation*.
- 9 In the Match Context area of the Constraints Table view, open the Resource drop down list, then select any resource to refresh the Constraints Table and Facts Table views.



The Policy Debugger displays a red icon near the constraints that fail to match. The larger, red octagonal icon shows the particular constraint that is “blocking” and preventing the job from running on the resource. This is the constraint that is causing the job to be in a “waiting” state. The Constraints Table also displays the policy name (debuggerExample) that is contributing the constraint that is causing problems.

There are a few ways to get the job to run:

- ♦ Create a Resource group called debugger, then place a resource in that group to satisfy the constraint specified in the policy.
- ♦ Disassociate the policy (debuggerExample) from the job (quickie).
- ♦ In the Constraints Table, right-click on the blocking constraint and select *Disable Constraint*.

Grid Object Health Monitoring

A

This section includes the following information:

- ♦ [Section A.1, “Health Facts,” on page 185](#)
- ♦ [Section A.2, “Health Events,” on page 187](#)
- ♦ [Section A.3, “The Health Debugger,” on page 187](#)

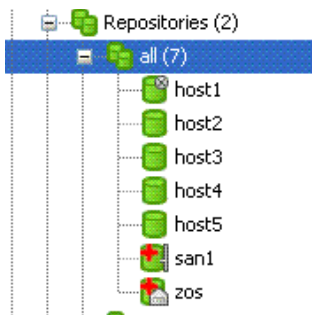
A.1 Health Facts

The Resource grid object, the VM host grid object, the User grid object and the Repository grid object each has an attribute or “Fact” that denotes the health of the object.

- ♦ `resource.health`
- ♦ `vmhost.health`
- ♦ `user.health`
- ♦ `repository.health`

Empirically, object health is a simple Boolean value, with `True` indicating that the object is healthy. This value can be controlled in a number of ways. An unhealthy object is displayed in the PlateSpin Orchestrate Development Client with a red cross to signal the object’s condition.

Figure A-1 Tree View of Repository Grid Objects in the “all” Group, Some Objects Unhealthy




You can define what constitutes the health or non-health of the grid object by setting this health fact. The health fact can be set or cleared in several ways:

- ♦ Explicitly set or cleared by the administrator using tools in the PlateSpin Orchestrate Development Client.
 - ♦ Select any grid object in the Development Client, then click the *Info/Groups* tab in the Workspace view. This is the “info” attribute editor. The attributes on this page let you edit facts.

The object information panel of the page has a *Healthy* check box that you can select or deselect to set the health of the object.


Figure A-2 The “Healthy” Check box on the Info/Groups Page



The screenshot shows a dialog box titled "User Information". Inside, there are three items: "Account Enabled" with a checked checkbox, "Online" with an unchecked checkbox, and "Healthy" with a checked checkbox.

- ♦ On the Constraint/Fact fact page of a grid object, right click the xxx.health fact name, then click *Edit/View Fact* to open the Edit Fact dialog box.

Figure A-3 The Edit Fact Dialog Box Displayed from the Constraint/Fact Page



The screenshot shows a dialog box titled "Edit Fact". At the top, it says "user.health" with a checked checkbox. Below that, it says "Boolean Fact XML Text (XML is read only)". The main area contains a text editor with the following XML code:

```
1 <fact name="user.health" value="true" type="Boolean" />
2
3
4
5
6
7
8
9
10
11
12
13
14
15
```

At the bottom, there are "OK" and "Cancel" buttons.

You can set the health of the object by selecting or deselecting the health check box. Changing the value in the Development Client in this way has an immediate effect unless the value is overridden by an attached policy (this follows the normal rules of policy inheritance).

- ♦ Set by using a discovery job (a job periodically scheduled to run on resources and to explicitly set the health fact, much like it sets other discovered facts). In this case, the discovery job performs a setFact (xxx.health) from JDL code.
- ♦ Set by using a policy. This method has little practical use except for locking the value immediately to override the setting (that is, the typical policy behavior) on the grid object:

```
<policy>
  <fact name="resource.health" value="true" type="boolean" />
</policy>
```

- ♦ Set by using a computed fact. This method can be used to monitor the health according to a computed value. One applied scenario for this method might be a computed fact that performs a statistical analysis of historical load data, perhaps provided by the Metrics facility.
- ♦ Set automatically by using a health constraint. This is the most practical use and is best illustrated with examples.

Example 1: Define resources as “unhealthy” if their 10 minute load average is greater than 5>

```

<policy>
  <constraint type="health">
    <lt fact="resource.metrics.loadaverage.history.10_min"
value="5.0" />
  </constraint>
</policy>

```

You could attach this policy directly to the Resource grid object or to a Resource group (more practical).

Example 2: Define a user as unhealthy if he or she has no money in their account.

```

<policy>
  <constraint type="health">
    <ge fact="user.account.balance" value="0" />
  </constraint>
</policy>

```

You could attach this policy directly to the User grid object, or to a User group (more practical).

You can aggregate (that is, group together with “and” or “or”) health constraints by using normal rules of policy aggregation.

By default, PlateSpin Control runs health constraints every 30 seconds. To alter this interval, you must contact Novell Support.

A.2 Health Events

Each time the value of a health fact changes, an event is generated. This event can be subscribed to by long-running Jobs (see “[Receiving Event Notifications in a Running Job](#)” in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference*) or the event can be used to trigger Jobs in the Job Scheduler (see “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.5 Development Client Reference*). The event names are different for each object type. They are listed in the following table.

Table A-1 Event Names for Grid Objects

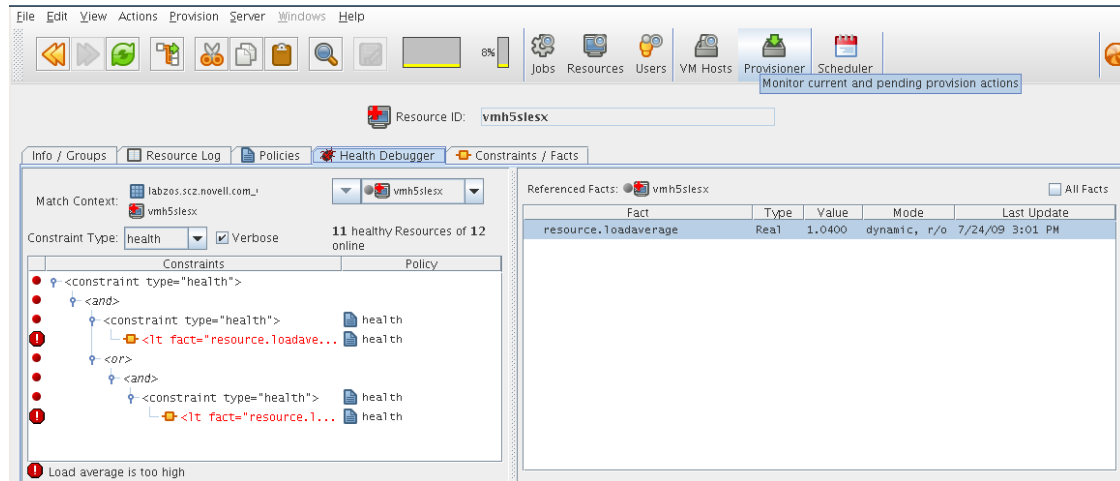
Object	Event Name
User	USER_HEALTH
Resource	RESOURCE_HEALTH
Repository	REPOSITORY_HEALTH
VmHost	VMHOST_HEALTH

A.3 The Health Debugger

Several objects modeled by PlateSpin Orchestrate have a health fact that can be used to visually show the health of the object or to trigger a job (see [Chapter 13, “The PlateSpin Orchestrate Job Scheduler,” on page 145](#)) upon a state change. This fact can optionally be manually set, or more usually automatically set through the periodic execution of the health constraint placed on that object. When such a health constraint is active, you might need to debug to discover the reasons for the changed state. The Health Debugger is a useful debugging tool.

The Health Debugger is a tabbed page (sometimes called an “admin view”) available in several views of the PlateSpin Orchestrate Development Client and functions much like the more generic “Policy Debugger.” This tool helps you to determine the reasons for the current Health of a Grid object. The following figure shows the Health Debugger tab opened in the Resource object view of the Orchestrate Development Client.

Figure A-4 PlateSpin Orchestrate Jobs View with the Health Debugger Page Open



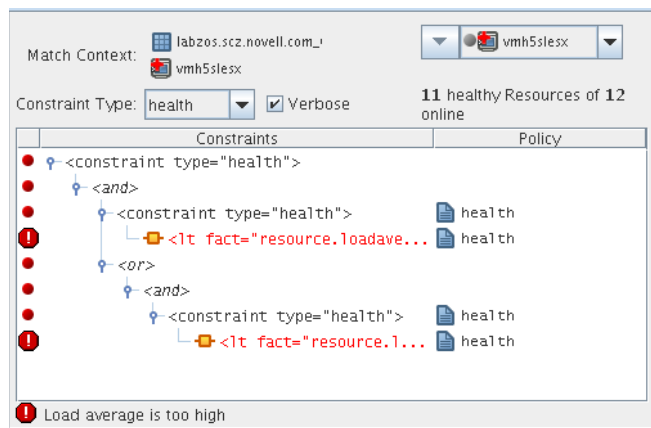
The Health Debugger tab is also available in the VM Host object view, the Repository object view and in the User object view of the Orchestrate Development Client.

- ♦ [Section A.3.1, “The Constraints Table Panel,” on page 188](#)
- ♦ [Section A.3.2, “The Facts Table View,” on page 191](#)

A.3.1 The Constraints Table Panel

The left side of the Health Debugger page is referred to as the Constraints Table view.

Figure A-5 The Constraints Table View



The appearance of this view can change, depending on the Constraint Type you select in the Constraint drop down list. In effect, if you change from the health selection, you will be debugging other constraints. For a description of these constraint types, see [Section 14.1.2, “The Constraint Type List,” on page 178](#). Different objects selected for the view change the *Match Context* in which the constraint is executed, which is useful for comparing how the constraint evaluates other objects.

The Constraints Table View is composed of several parts:

- ♦ “The Match Context Area” on page 189
- ♦ “The Verbose Check Box” on page 190
- ♦ “The Capable Objects Summary” on page 190
- ♦ “The Constraints Column of the Constraints Table View” on page 190

The Match Context Area

The Health Debugger provides the general identification of a Grid object in the *Match Context* area of the Constraints Table View. The Match Context defines every object that is available to the constraint you are viewing.

Figure A-6 The Match Context Area of the Constraints Table View in the Health Debugger



In this example, the identifying Facts in the Match Context include:

- ♦ Matrix: The icon and a text string identifies the machine that matches the grid name given to the Orchestrate Server where this object exists.
- ♦ Object icon: The object icon and a text string identifies the object (VM host, Repository, or that matches the user who is running the job. If the object icon has a red cross overlaid, it is unhealthy.
- ♦ Object list: The object drop down list shows all named objects of the type selected in the Explorer Tree. The objects in the list that are currently offline display with a dimmed icon. If available, a listed object has a colored dot by its side. The color of the dot (blue ● or gray ●) and the object type it accompanies has significance:
 - ♦ A blue dot with the *All <Object Type>* label indicates that at least one object in the list matches the constraints and is healthy.
 - ♦ A gray dot with the *All <Object Type>* label indicates that no objects of this type match the constraints.
 - ♦ A blue dot with a named, selected object indicates that its constraints match and it is healthy.
 - ♦ A gray dot by a named, selected object indicates that its constraints do not match and that it is not healthy.

The Verbose Check Box

When you select the Verbose check box, the reason string specified in the policy is displayed in the *Constraints* tree. For more information, see [Section 14.1.5, “The Constraints Column of the Constraints Table View,”](#) on page 178.

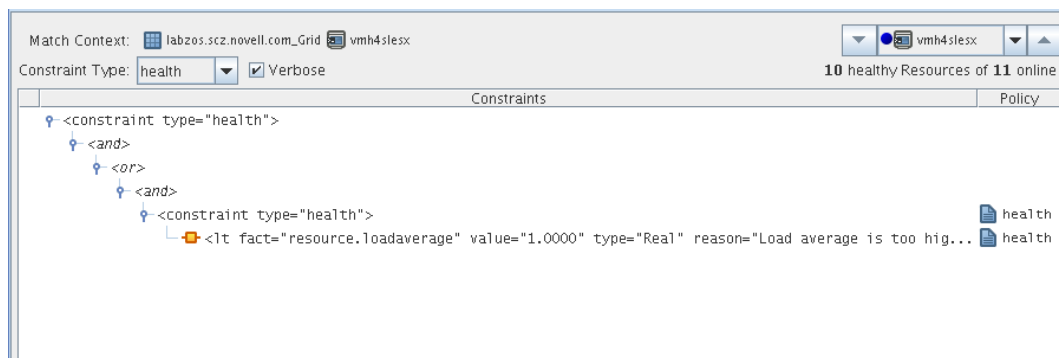
The Capable Objects Summary

Directly under the Object list in the constraint view, a populated string summarizes the resources that are capable of servicing the job. For example, 11 healthy Resources of 12 online indicates that 11 of the 12 total online Resources are healthy.


The Constraints Column of the Constraints Table View

The Constraints column of the constraints table view shows the logical hierarchy (that is, a “tree” format) of the constraints that are defined by the Policies associated with the Job. You can identify the status of the listed constraints by the icons that might be displayed in the far left column of the table:

- ♦ no icon: The constraint passes the match. It is a “true” match. No red icons are displayed next to any listed constraint.

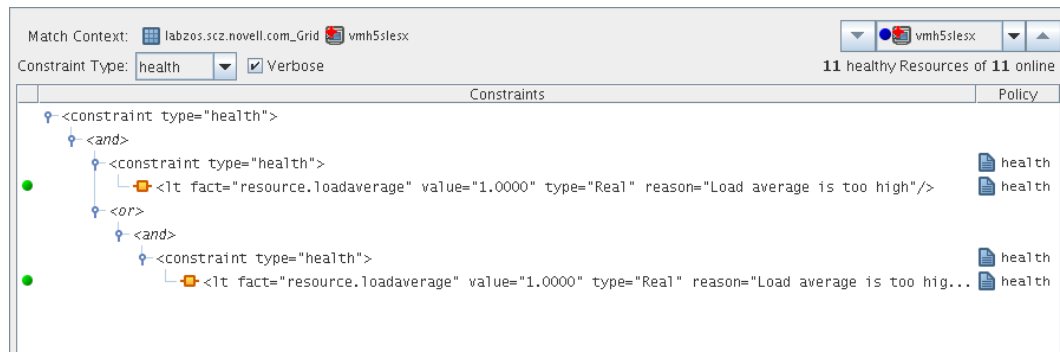


- ♦ red dot icon: The constraint does not pass the match. The figure below shows that the resource vmh5slesx cannot run the job because its health constraints do not match.



- ♦ red octagonal icon: For convenience, just one of the constraints is identified as the blocking constraint. This is the constraint that the system has determined as responsible for the constraint as a whole to fail (note that individual constraint lines can fail without causing the the entire constraint to fail).

- ♦ green dot icon: A failing constraint has been disabled so that it behaves like a match (pass). The figure below shows the green dot icon next to that the constraint that was formerly failing and can now be forced to behave as a match.




If you right-click a constraint in the table, a popup menu with three options is displayed. This menu lets you change the status of the constraint. Disabling a constraint is useful if you want to temporarily relax a condition without editing or redeploying the whole policy and potentially affecting other objects that share the policy. A disabled constraint can be re-enabled later.

- ♦ **Show Admin View:** Select this option to open the Admin View for the specific object selected.
- ♦ **Disable Constraint:** Select this option to disable (attach a green dot icon to) a constraint. Disabling a constraint with this function effectively makes it match, a condition that can prove useful if you want to perform a “what if” test without actually changing a policy.
- ♦ **Enable All Constraints:** Select this option if you have disabled one or more constraints during testing and you want to restore them to the enabled state.

NOTE: Health constraints are always re-evaluated in the debugger. The last system execution (cached constraint) is not available for health constraints.

The Policy column of the constraints table displays the policy name that contributed the constraint. Right-click a policy name to open a popup menu offering the option to open the policy editor for the specified policy. The menu also includes constraint enabling or disabling options, just as the popup menu for constraint column does.

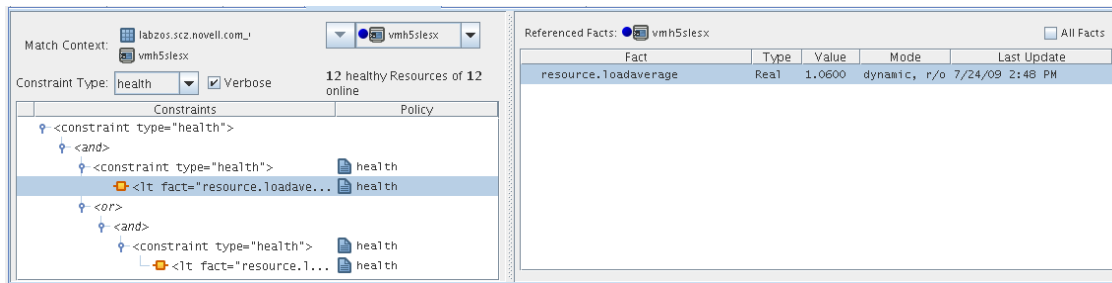
Figure A-7 The Popup Menu Launched from the Policy Column



A.3.2 The Facts Table View


The Facts Table view displays the facts referenced in the Constraint Tree view for a specified object. Selecting a fact in the Constraint tree automatically selects that fact in the table.

Figure A-8 The Constraints Table View and the Accompanying Facts Table View



If you right-click a column head in this table, a menu is launched where you can select the columns that you want to display.

Figure A-9 Menu Used to Select the Columns Displayed in the Facts Table View of the Policy Debugger



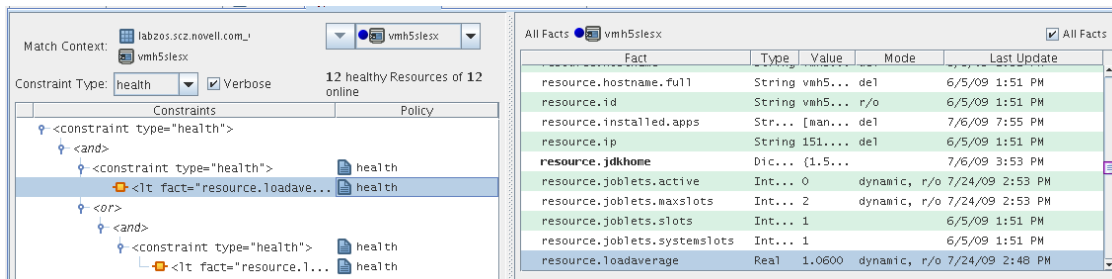
- ◆ “The All Facts Check Box” on page 192

The All Facts Check Box

If you select the *All Facts* check box at the top of the Facts Table view, all of the facts (including matrix, and <object type> facts) associated with the Match Context are listed.

If you select *All <Object Type>* in the Match Context (see [Section 14.1.1, “The Match Context Area,” on page 176](#)) and you also select the *All Facts* check box, the Facts Table view displays all the <object type> facts for the specified Match Context of the selected object.

Figure A-10 All Facts Check Box Selected with All VM Hosts in Match Context




This section contains the following information:

- ♦ [Section B.1, “Event Object Visualization and Management in the Development Client,” on page 193](#)
- ♦ [Section B.2, “The Event Debugger,” on page 196](#)
- ♦ [Section B.3, “Understanding the PlateSpin Orchestrate Events System,” on page 200](#)

B.1 Event Object Visualization and Management in the Development Client

The Events folder is displayed in the Explorer tree between the Computed Facts and Metrics folders.


Figure B-1 *The Events Folder in the PlateSpin Orchestrate Explorer View*




Although the PlateSpin Orchestrate system includes several built-in Events (see [Section B.3.2, “Built-in Events,” on page 201](#)), these Events are not displayed in the Explorer view. Only custom Events (defined in XML by the administrator and then deployed on the server) are displayed in the tree.

When an Event object is deployed, its icon is displayed in the tree in the Events folder.

Figure B-2 *An Event Object in the Events Folder*



The icon in the Explorer tree might be overlaid with a write symbol  to indicate that its XML content has changed and needs to be saved. For more information about changing the XML content, see [Section B.1.4, “The Event Editor,” on page 195](#).


This section includes the following information about how the Event object is managed in the Development Client:

- ♦ [Section B.1.1, “Deploying a New Rule-Based Event,” on page 194](#)
- ♦ [Section B.1.2, “Deploying a Pre-written Rule-Based Event,” on page 194](#)
- ♦ [Section B.1.3, “Undeploying an Event,” on page 195](#)
- ♦ [Section B.1.4, “The Event Editor,” on page 195](#)

B.1.1 Deploying a New Rule-Based Event

Use the following steps to create a new event.

- 1 In the Explorer view, right-click the Events folder, then select *New Event* to open the Create a New Event dialog box.




- 2 Enter the name for the new Event, then click *OK* to create the new Event object.

PlateSpin Orchestrate then deploys the new Event object on the server, where it can be managed. The Development Client opens the Event Editor, where you can edit the XML definition of this Event. For more information, see [Section B.1.4, “The Event Editor,”](#) on page 195 and [Section B.3.3, “Rule-based Events,”](#) on page 202.

B.1.2 Deploying a Pre-written Rule-Based Event

Use the following steps to deploy a pre-written Event (an XML .event file).

- 1 Right-click the Events container, then select *Deploy Event* to open the Select the Component File to Deploy dialog box.



- 2 In the dialog box, navigate to the file system location of the Event file you previously created, or to an example .event file from /opt/novell/zenworks/zos/server/examples/events, then click *OK* to deploy the pre-written Event.

When you deploy the rule-based Event, the Development Client opens the Event Editor, where you can edit the XML definition of this Event. For more information, see [Section B.1.4, “The Event Editor,” on page 195](#) and [Section B.3.3, “Rule-based Events,” on page 202](#).

B.1.3 Undeploying an Event


When an Event has been deployed, it can be undeployed. Undeploying deletes the Event object within PlateSpin Orchestrate, but it does not delete the source .event file, which still exists and can be redeployed.

To undeploy an Event, right-click on the Event object in the Explorer tree, then select *Undeploy*. You can also simply select the object and press Delete.

B.1.4 The Event Editor


The Event Editor opens when you select a deployed Event in the Explorer tree of the Development Client.

Figure B-3 The Event Editor



The screenshot shows a software window titled "Event Editor" with a sub-tab "Event Debugger". The main area contains XML code for an event, with line numbers 1 through 45 on the left. The code includes comments explaining the event's purpose and behavior, and XML elements for context, trigger, and reset rules.

```
1  <!--
2    Example event that triggers when a Vm Host becomes overloaded.
3
4    The context for this event includes all Vm Host's in the system. This is
5    specified by the <vmhost/> XML element. This means that the
6    <trigger> rule is evaluated on every Vm Host.
7
8    In this case the Trigger rule evaluates whether the metric for a historical
9    load average exceeds a specified threshold.
10
11    To force this trigger to pass, you can manually increase the load on a
12    physical vmhost or lower the threshold value in the <gt> rule.
13
14    Once the the trigger rule passes then any defined event notifications
15    such as starting a Scheduled Job and/or notify a long running job will
16    occur. For example, you can use the Job Scheduler to start a job that
17    migrates Vms from the loaded Vm Host that triggered the rule.
18    This is shown in the migrate.jdl example job.
19
20    The trigger will pass again if a different Vm Host passes the rule.
21
22    The same Vm Host will no longer pass the trigger until the reset rule
23    passes.
24  -->
25
26  <event>
27
28    <context>
29
30      <vmhost />
31
32    </context>
33
34
35    <trigger>
36      <gt fact="vmhost.vm.count" value="0" />
37      <gt fact="vmhost.resource.metrics.load_one.raverage" value="2" />
38    </trigger>
39
40
41    <reset>
42      <lt fact="vmhost.resource.metrics.load_one.raverage" value=".5" />
43    </reset>
44
45  </event>
```

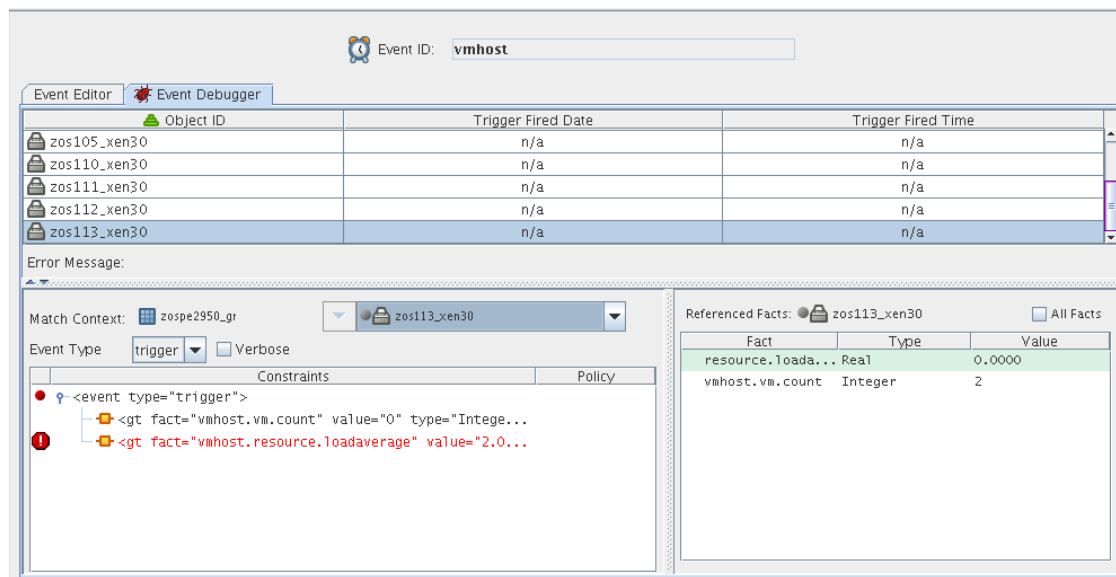
Inside this editor, you can make changes to the XML content of the Event. Example Events contain comments that explain how you can use them and the behavior you can expect to see as a result of deploying them. For the changes you make to be effective, you need to click the Save  tool.

For more information about the allowed XML syntax within an Event, see [Section B.3.3, “Rule-based Events,”](#) on page 202.

B.2 The Event Debugger

The Event Debugger is a tabbed page available from the Explorer view of an Event object when you select an event and then click the Event Debugger tab. This tool helps you to determine the reasons for the current state (triggered or reset) of an Event. The following figure shows the Event Debugger view.

Figure B-4 *The Event Debugger*




The information in this section describes the various parts of the debugger.

- ♦ [Section B.2.1, “The Constraints Table,” on page 197](#)
- ♦ [Section B.2.2, “The Facts Table,” on page 199](#)

B.2.1 The Constraints Table

The left side of the Event Debugger panel includes the Constraints Table. The rules that define the `<trigger>` and `<reset>` of an Event are defined using the same XML constraint syntax used in policies.

Figure B-5 *The Constraints Table Area of the Events Debugger*




The Constraints Table has several parts:

- ♦ “Match Context” on page 198
- ♦ “Event Type List” on page 198
- ♦ “Verbose Check Box” on page 198
- ♦ “Constraints List (Tree)” on page 198

Match Context

Depending on the Event, the debugger identifies the Grid objects (Job, Jobinstance, Resource, Repository, User, or VMHost) that define the context of the trigger or reset rules specified in the Event XML.

Figure B-6 The Match Context Area of the Constraints Table View in the Event Debugger



Event Type List

Select one of the Event types in the drop down list to debug how either the <trigger> or <reset> rule(s) are being applied. Constraint types in the list are disabled (dimmed) if they do not apply to the Event that you are debugging.

- ♦ **trigger:** The rules defining the conditions (through a constraint) in which an Event is generated.
- ♦ **reset:** The rules defining the conditions (through a constraint) in which an Event is reset (that is, able to be triggered again).

Verbose Check Box

When you select the Verbose check box, additional constraint information is displayed.

Constraints List (Tree)

The Constraints Tree, a column in the constraints table, lists the constraints that constitute a particular rule in a hierarchical view.

Each constraint is flagged with an icon to signify whether it “passes” or not. A constraint flagged with an exclamation point indicates a constraint causing the rule to not “pass.”

Right-click on a constraint to expose a menu where you can perform one of the following actions:

Show Admin View: Selects the currently evaluated Grid object in the Explorer Tree and displays its Info/Groups administration information view.

Disable Constraint: “Passes” the constraint, regardless of how it evaluates.


Enable All Constraints: Re-enables any disabled constraints.

NOTE: The right-click menu is available only when you select specific constraints.

B.2.2 The Facts Table

The Facts Table view displays the facts referenced in the Constraint Tree view for a specified Event. Selecting a rule containing a particular fact in the Constraint tree automatically selects that fact (and its current value) in the table.

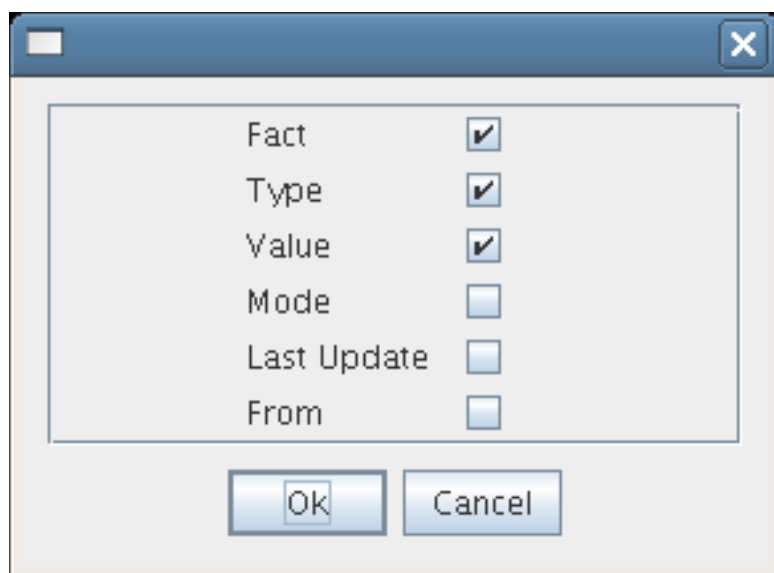
Figure B-7 The Facts Table View



Fact	Type	Value
resource.10a...	Real	0.0000
vmhost.vm.count	Integer	4

If you right-click a column head in this table, a menu is launched where you can select the columns that you want to display.

Figure B-8 Menu Used to Select the Columns Displayed in the Facts Table View of the Policy Debugger



All Facts Check Box: If you select the *All Facts* check box at the top of the Facts Table view, all facts for the selected Grid object, as well as those for the Server itself (`matrix.* facts`) are displayed.

If you right-click a fact, you have the option of adding a new fact, deleting the selected fact, or viewing/editing (if the fact is editable or can be deleted) the selected fact.

B.3 Understanding the PlateSpin Orchestrate Events System

The PlateSpin Orchestrate Event System integrates with the Job Scheduler. Event notifications can start jobs and can also invoke Event handler methods in long-running jobs. In turn, a job can react to the Event by starting other PlateSpin Orchestrate actions, by modifying object attributes, or by executing another external process.

For example, an Event notification can occur when a VM Host has exceeded its configured load limits. This Event can start a job that migrates VMs off of the loaded VM Host or VM Hosts.

PlateSpin Orchestrate supports two Event types:

- ◆ Built-in Events, such as change of status of the health of a resource's or a change in the online status of the resource.
- ◆ Rule-based Events that are triggered when the attributes of an object satisfy the rules (constraint conditions) defining the Event.

This section includes the following information:

- ◆ [Section B.3.1, “Event Notification,” on page 201](#)
- ◆ [Section B.3.2, “Built-in Events,” on page 201](#)
- ◆ [Section B.3.3, “Rule-based Events,” on page 202](#)

B.3.1 Event Notification

An Event notifies two other PlateSpin Orchestrate services— the Job Scheduler and the Job Broker. The Job Scheduler starts jobs that are awaiting an Event to trigger them.

The Job Broker invokes a callback on any long-running job that has registered for notification of an Event.

See “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.5 Development Client Reference* for more information about setting up a Job Schedule. Also see “[Job Architecture](#)” in the *PlateSpin Orchestrate 2.5 Developer Guide and Reference* for more information about how jobs can register for event notification.

B.3.2 Built-in Events

Built-in events occur when a managed object comes online/offline or has a health status change.

PlateSpin Orchestrate uses the following built-in Events to keep managed objects synchronized.

Table B-1 *PlateSpin Orchestrate Built-in Events*

Event Name	Description
AGENT_VERSION_MISMATCH	Resource Agent version mismatch (agent needs upgrade)
REPOSITORY_HEALTH	Repository health status has changed
RESOURCE_HEALTH	Resource health status has changed
RESOURCE OFFLINE	Resource Agent has logged off from server
RESOURCE_ONLINE	Resource Agent has logged in to server
SERVER_UP	Server has fully started
USER_HEALTH	User health status has changed
USER_ONLINE	User has logged in to server
VMHOST_ADDED	VM Host has been added
VMHOST_HEALTH	VM Host health status has changed
VMHOST_NOT_AVAILABLE	No VM Host is available
VMHOST_REMOVED	VM Host has been removed

For example, when a resource comes online (that is, the agent connects to server), the `RESOURCE_ONLINE` Event is fired and both Scheduled Jobs with a trigger for that Event and long-running jobs with Event handlers are notified.

The `RESOURCE_ONLINE` built-in Event is used by the embedded discovery jobs, such as for discovering operating system and CPU information (`osInfo` and `cpuInfo` jobs). Both `osInfo` and `cpuInfo` job archives (`.job`) include a schedule file (`.sched`) specifying a trigger (`.trig`) that allows these jobs be started when notification of the `RESOURCE_ONLINE` Event occurs.

B.3.3 Rule-based Events

Rule-based Events are defined in an XML document and deployed to the PlateSpin Orchestrate Server and managed using the PlateSpin Orchestrate Development Client. Rules can be a simple object attribute (fact) equivalency check or they can use AND,OR, IF, ELSE logic, among other things, in an Event ruleset.

The rules follow the same syntax as the constraints that are defined in XML policy files for all Grid Objects, such as Jobs, VM Hosts, etc.

The PlateSpin Orchestrate Event Service evaluates the rules; if the rules pass, an Event notification occurs.

The XML Schema document specification can be found in `<install dir>/doc/xsds/event_1_0_0.xsd`.

The Event XML specification is composed of three sections.

- ♦ `<context>`
- ♦ `<trigger>`
- ♦ `<reset>`

NOTE: Both the `<context>` and `<trigger>` sections are required.

`<context>` section

The context section defines the context in which the Event rules are evaluated. With Events, you specify what objects are in the Event rule context in this section. The available objects are `Job`, `Jobinstance`, `Resource`, `Repository`, `User`, and `VMHost`. From these objects, you can specify one object set to iterate over and optionally a single instance of the object.

`<trigger>` section

The trigger section defines the rules for when an Event notification occurs. The `<trigger>` format is the same syntax as `<constraints>` used in policies.

`<reset>` section

The optional `<reset>` section defines the rules for when an Event is reset. If the `<reset>` rule is not used, an Event is reset based on a timeout. The `<reset>` format is also the same syntax as in `<constraints>` used in policies.

The `resetInterval` attribute is set on the `<event>` XML element. If "resetInterval" and `<reset>` is not used, the default timeout for resetting is 10 minutes.

The following example (taken from the "vmhost.event" in `<install dir>/examples/events`) defines that a notification occurs when a VM Host becomes overloaded.

```

1<event>
2
3  <context>
4    <vmhost />
5  </context>
6
7  <trigger>
8    <gt fact="vmhost.vm.count" value="0" />
9    <gt fact="vmhost.resource.loadaverage" value="2" />
10 </trigger>
11
12 <reset>
13   <lt fact="vmhost.resource.loadaverage" value=".5" />
14 </reset>
15
16</event>

```

Lines 3-5: Defines the context for the Event's rule evaluation.

Line 4: The context specifies for all VM Host objects, so the Event Service iterates over all VM Hosts. On each VM Host, the `<trigger>` rule will be evaluated, so in this case, the Event context is composed of one or more VM Hosts.

Lines 7-12: Defines the Trigger rule to determine if this Event is to fire notifications or not. If the trigger rule does not pass, then no Event notifications occur.

Line 8: Consider only VM Hosts that have at least one VM instance running.

Line 9: Check the running average of the VM Host's load average if it exceeds a threshold value. In this case if the average is greater than 2.

Lines 12-14: Defines the Reset rule to determine if a previously triggered VM Host can be reset and triggered again.

Line 13: Only reset if the running average of the VM Host's load average drops below a threshold.

So when a VM Host passes the trigger rule, the VM Host does not pass the trigger rule again until the reset rule (load average drops below threshold) passes.

See the `repository.event` example (`<install dir>/examples/events/repository.event`) for an Event with a rule that evaluates the `freespace` fact on all repository objects.

The Metrics Facility

C

The Orchestrate Metrics Facility collects, aggregates, and allows simple fact-based retrieval of metric values by jobs and computed facts (via JDL), policy constraints, and Event triggers on a per-resource basis. This provides aggregated metrics generated by gmond without the need for the gmetad Ganglia service. Note that gmetad can still be used in parallel for aggregating gmond reported metrics for visualization purposes.

- ♦ [Section C.1, “Metrics Facility Functionality,” on page 205](#)
- ♦ [Section C.2, “Ganglia Metrics,” on page 205](#)
- ♦ [Section C.3, “How Does the Metrics Facility Impact Orchestrate Server Performance?,” on page 207](#)
- ♦ [Section C.4, “RRD Definition Using Deployable .metric Files,” on page 207](#)
- ♦ [Section C.5, “Query of Aggregated Metric Values,” on page 209](#)
- ♦ [Section C.6, “MetricsManager MBean API,” on page 210](#)
- ♦ [Section C.7, “Using the Metrics Facility in the Development Client Interface,” on page 211](#)

C.1 Metrics Facility Functionality

The Metrics Facility provides the following functionality:

- ♦ Collection of gmond provided metrics using the Orchestrate Agent .
- ♦ Retrieval of instantaneous metric values via `resource.metrics.<METRIC_NAME>` fact space, where `<METRIC_NAME>` is the name of the metric.
- ♦ Deployable Round Robin Database (RRD) (data aggregation) definition using XML `.metric` files, which allows flexible definition of aggregation periods. For example, using `resource.metrics.<METRIC_NAME>.10_minute.average` as a 10-minute aggregation period separate from `resource.metrics.<METRIC_NAME>.1_hour.average`.
- ♦ Retrieval of an array of aggregated metric values using `resource.metrics.<METRIC_NAME>.xxx.values`.
- ♦ Zero-configuration for core Ganglia metrics. The Orchestrate Agent automatically discovers if gmond is running on a resource, and the Orchestrate Server collects and, if a `.metric` file is configured, aggregates those metrics.
- ♦ Persistence of collected RRD data across server restart and high availability fail-over conditions.

C.2 Ganglia Metrics

Not all Ganglia metrics are suitable for aggregation, such as those of “String” type. By default, only the 24 metrics listed in the following table are supported (either of type “Real” or type “Integer”):

Table C-1 *Supported Ganglia Metrics*

Metric	Type
bytes_in	Real
bytes_out	Real
cpu_idle	Real
cpu_nice	Real
cpu_system	Real
cpu_user	Real
cpu_wio	Real
disk_free	Integer
disk_total	Integer
load_fifteen	Real
load_five	Real
load_one	Real
mem_buffers	Integer
mem_cached	Integer
mem_free	Integer
mem_shared	Integer
mem_total	Integer
pkts_in	Real
pkts_out	Real
proc_run	Integer
proc_total	Integer
swap_free	Integer
swap_total	Integer

C.3 How Does the Metrics Facility Impact Orchestrate Server Performance?

The Metrics Facility balances flexibility and minimizing the impact on overall Orchestrate Server performance. There is no continuous parsing of XML on the server. Instead, parsing of gmond-generated XML is performed by each managed resource.

This section includes the following information:

- ♦ [Section C.3.1, “I/O Contention,” on page 207](#)
- ♦ [Section C.3.2, “Too Many Open Files,” on page 207](#)

C.3.1 I/O Contention

By default, RRD-based data aggregation is file-based. Because of the frequency of updates and queries of RRD files, this poses a significant performance issue. The Metrics Facility minimizes I/O contention by using in-memory caching and batched write operations to avoid I/O contention and the resulting performance degradation.

C.3.2 Too Many Open Files

The PlateSpin Orchestrate Monitoring Server uses a “one RRD per resource” approach, where a RRD contains the AVERAGE, MIN, and MAX RRAs for multiple metrics (DS). In contrast, the Metrics Facility takes an “inside out” approach, which results in “one RRD per metric.” For 1,000 agents reporting 24 Ganglia metrics, this reduces the number of files dramatically (from 24,000 in the “one RRD per metric, per resource” case, and 1,000 in the “one RRD per resource” case, to 24 in the “one RRD per metric” case). This approach avoids a “too many open files” condition.

NOTE: A “too many open files” condition occurs when the default maximum file descriptors available to a process launched from the Linux shell is exceeded.

C.4 RRD Definition Using Deployable .metric Files

Definition of the “aggregation” functions performed by the Metrics Facility’s internal RRD data structures are customizable using deployable XML .metric definition files. This accommodates a flexible configuration of the following:

- ♦ The Ganglia-reported metrics to include in aggregated data structures
- ♦ The data aggregation periods that are of specific interest

The deployable definition files, one per metric to be aggregated, consist of the following:

- ♦ The name of the “instantaneous” metric to be aggregated, for example, `load_one`
- ♦ An optional description of the metric to be aggregated.
- ♦ A “heartbeat” value that governs the updates to the contained RRAs.
- ♦ One or more named Periods (corresponding to the RRAs to be created) with an optional description.

- ♦ For each Period, the number of data points to aggregate (steps), the number of aggregated data points to archive (rows), and the “xff” (x-files factor) allowed for the metric.

NOTE: The xff determines how many of the samples can be NaN for the consolidated sample to be considered NaN. Usually, this is set to 0.5, or 50%.).

When an RRD is defined through deployment of its definition file, three RRAs are created for each Period: AVERAGE, MAX, and MIN. A new DS (datasource) is added to the RRD for each resource reporting the metric to be aggregated. This requires the RRD file to be re-created each time a new resource begins reporting a given metric and the previously aggregated values copied from the old RRD to the new one. This approach enhances performance and flexibility, but the RRD file is not of fixed size: Over time, the RRD grows or shrinks as new resources are added to the system or are deleted from the Orchestrate model.

NOTE: The RRD is actually re-created with a new DS added for each new resource and the “old” RRA’s data copied into it.

Deleting a Resource GridObject removes its DS from the RRD file (actually, from all RRDs with metrics reported by that resource).

One optimization you can implement for storing the smallest Period (consisting of a single step) is to create only a single RRA (vs. three), because the average of a single datapoint is equal to the maximum and minimum of a single datapoint.

NOTE: The RRD files created by the Java rrd4j library are not binary compatible with RRD files generated by the rrdtool used by gmetad. They are however portable across operating system architectures (e.g., 32-bit bigendian vs. 64-bit little-endian) which is not possible with traditional RRD files created using rrdtool.

C.4.1 XML Format for Deployable .metric Definitions

An example of the format of the deployable RRD definition is shown below.

```
<metric name="load_one" heartbeat="120"
    description="Ganglia oneminute load average">

    <period name="1_minute" steps="1" rows="60" xff=".5"
        description="1 hour worth of 1 minute (raw) data"/>

    <period name="5_minute" steps="5" rows="12" xff=".5"
        description="1 hour worth of 5 minute aggregations"/>

    <period name="10_minute" steps="10" rows="72" xff=".5"
        description="12 hours worth of 10 minute aggregations"/>

    <period name="1_hour" steps="60" rows="24" xff=".5"
        description="1 days worth of 1 hour aggregations"/>

</metric>
```


This example creates an RRD for the `load_one` metric, with four aggregation periods (RRAs) called `1_minute`, `5_minute`, `10_minute`, and `1_hour`. The default sample (RRD update) time is one minute, so the `10_minute` aggregation period has 10 steps. 72 “rows” of aggregated datapoints are retained in the RRA before the oldest is dropped off, representing 12 hours ($12 * 60 / 10$) worth of data.

NOTE: The `1_minute` period is not a true aggregation because the default sample (RRD update) time is also one minute. In this case, the “raw” datapoints are stored for historical reference.

C.5 Query of Aggregated Metric Values

Aggregated metric values can be queried similar to the instantaneous values either from JDL, from within a policy/event constraint, or from an array of multiple metrics values. This section includes the following information:

- ♦ [Section C.5.1, “Example of a JDL Query for Aggregated Metric Values,” on page 209](#)
- ♦ [Section C.5.2, “Example of a Policy Constraint or Event Constraint Using Aggregated Metric Values,” on page 209](#)
- ♦ [Section C.5.3, “Example of Using Non-aggregated \(“Raw”\) Historical Metric Values,” on page 210](#)

C.5.1 Example of a JDL Query for Aggregated Metric Values

```
r = getMatrix().getResource("local_vmhost")

print "resource.id: %s" % (r.getFact("resource.id"))

print "load_one.10_minute.average: %s" %
( r.getFact("resource.metrics.load_one.10_minute.average") )
```

C.5.2 Example of a Policy Constraint or Event Constraint Using Aggregated Metric Values

```
<event>
  <context>
    <vmhost />
  </context>
  <trigger>
    <gt fact="vmhost.vm.count" value="0" />
    <gt fact="vmhost.resource.metrics.load_one.10_minute.average"
value="2" />
  </trigger>
  <reset>
    <lt fact="vmhost.resource.metrics.load_one.10_minute.average"
value=".5" />
  </reset>
</event>
```

The consolidation functions (AVERAGE, MAX, and MIN) are supported for each defined aggregation period. In RRD terminology, this means that for each metric, there are three RRAs defined for each “period” element in the `.metric` deployable definition.

C.5.3 Example of Using Non-aggregated (“Raw”) Historical Metric Values

You can query for an array of “raw” values that constitute the aggregated datapoints for a given RRA within the RRD data structure by appending `.values` to the factname representing the metric (and period) of interest. For example, to print all the MIN values collected for the `10_minute` aggregation period, the JDL is

```
r = getMatrix().getResource("local_vmhost")

print "resource.id: %s" % (r.getFact("resource.id"))
print "load_one.10_minute.min.values: %s" %

( r.getFact("resource.metrics.load_one.10_minute.min.values") )
```

Similarly, to print the array of AVERAGE and MAX values collected, the JDL is

```
print "load_one.10_minute.average.values: %s" %
( r.getFact("resource.metrics.load_one.10_minute.average.values") )

print "load_one.10_minute.max.values: %s" %
( r.getFact("resource.metrics.load_one.10_minute.max.values") )
```

C.6 MetricsManager MBean API

The MetricsManager facility exposes a small number of methods for disabling and enabling a RRD file creation and update, and for query of both instantaneous and aggregated metric values. The facility also allows a UI to query for information needed for populating a pull-down list that could include, for example, valid metric names for a specified resource or valid aggregation periods for such a metric. The API also currently provides a way to fetch a “running average” of the raw datapoints for a metric (a cached value that does not require a fetch operation from an RRD) and server-side generation of a simple graph of metric data (requires a fetch from an RRD file).

This section includes the following information:

- ♦ [Section C.6.1, “MBean Methods Exposed by the MetricsManager Facility,” on page 210](#)
- ♦ [Section C.6.2, “The MetricsDeployer Facility,” on page 211](#)

C.6.1 MBean Methods Exposed by the MetricsManager Facility

The following table lists the MBean methods that are exposed by the MetricsManager Facility.

Table C-2 MBean Methods Exposed by the MetricsManager Facility

MBean Method	Description
<code>setMetricsEnabled()</code>	A parameter of false disables the creation or update of RRD files.
<code>getObjectNames()</code>	Returns a list of current deployed metric definitions
<code>update()</code>	Updates RRDs with instantaneous values collected by the Metrics Facility.
<code>getFact()</code>	Given the fully qualified fact name, fetches the value of a specified resource’s fact.

MBean Method	Description
<code>getMetricNames()</code>	Fetches the metric names relevant for a specified resource .
<code>getPeriodNames()</code>	Fetches the names of periods (RRAs) defined for a specified metric.
<code>getRunningAverage()</code>	Fetches the running average of a metric for a specified resource.
<code>getRrdGraphic()</code>	Used by the Development Client to fetch a server-generated ImageIconInfo object.

C.6.2 The MetricsDeployer Facility

The MetricsDeployer deployment facility parses the `.metric` definition files and creates the associated grid objects that are used to maintain the metadata related to the RRD used for aggregation.

C.7 Using the Metrics Facility in the Development Client Interface

The PlateSpin Orchestrate Development Client supports the Metrics Facility in the following ways:

- ♦ **Metrics object:** A Metrics object is deployed in the Explorer tree. Use the right-click menu to display the “deploy” and “undeploy” actions. Pre-defined `.metric` files are located in the `/opt/novell/zenworks/zos/server/components/metrics` folder, or you can create a new `.metric` file and paste in the XML as shown in [Section C.4.1, “XML Format for Deployable .metric Definitions,” on page 208](#).

Metrics objects are listed by their deployment name, which may or may not be the same as the name of the actual metric. This potentially allows multiple, separately deployable, RRD definitions for a single “instantaneous” metric, with different aggregation periods defined.

- ♦ **Metrics Editor:** Selecting the Metrics object opens a free-form XML editor in the admin view of the client. The editor is similar to the Event” editor for viewing RRD definition XML.

Documentation Updates

D

This section contains information about documentation content changes that were made in this *PlateSpin Orchestrate Development Client Reference* after the initial release of PlateSpin Orchestrate 2.5. The changes are listed according to the date they were published.

The documentation for this product is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the changes listed in this section.

If you need to know whether a copy of the PDF documentation that you are using is the most recent, the PDF document includes a publication date on the title page.

The documentation was updated on the following dates:

- ♦ [Section D.1, “November 29, 2010,” on page 213](#)

D.1 November 29, 2010

Updates were made to the following sections:

Location	Update Description
Table 3-1 on page 33	Added detail about additional log channels and when messages are sent through them.
Section 12.4, “Metrics Objects,” on page 144	Added new content.
Appendix A, “Grid Object Health Monitoring,” on page 185	Consolidated the former <i>Appendix D</i> with this appendix. The content is now included in Section A.3, “The Health Debugger,” on page 187 .
Former <i>Appendix C, “Provisioning Actions and History.”</i>	Moved this content to “ Provisioning Actions and History ” in the <i>PlateSpin Orchestrate 2.5 Virtual Machine Management Guide</i> .
Appendix C, “The Metrics Facility,” on page 205	Added new content.

