

Developer Guide and Reference

Novell. PlateSpin. Orchestrate

2.0.2

July 9, 2009

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 © 2008-2009 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

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

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

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	11
1 Getting Started With Development	15
1.1 What You Should Know	15
1.1.1 Prerequisite Knowledge	15
1.1.2 Setting Up Your Development Environment	16
1.2 Prerequisites for the Development Environment	16
2 Advanced Job Development Concepts	19
2.1 JDL Job Scripts	19
2.1.1 Principles of Job Operation	20
2.2 Understanding TLS Encryption	21
2.3 Understanding Job Examples	21
2.3.1 provisionBuildTestResource.job	21
2.3.2 Workflow Job Example	23
3 The PlateSpin Orchestrate Datagrid	25
3.1 Defining the Datagrid	25
3.1.1 PlateSpin Orchestrate Datagrid Filepaths	25
3.1.2 Distributing Files	26
3.1.3 Simultaneous Multicasting to Multiple Receivers	26
3.1.4 PlateSpin Orchestrate Datagrid Commands	27
3.2 Datagrid Communications	27
3.2.1 Multicast Example	28
3.2.2 Grid Performance Factors	28
3.2.3 Plan for Datagrid Expansion	29
3.3 datagrid.copy Example	29
4 Using PlateSpin Orchestrate Jobs	31
4.1 Resource Discovery	31
4.1.1 Provisioning Jobs	31
4.1.2 Resource Discovery Jobs	31
4.2 Resource Selection	32
4.3 Workload Management	33
4.4 Policy Management	34
4.5 Auditing and Accounting Jobs	35
4.6 BuildTest Job Examples	35
4.6.1 buildTest.policy Example	36
4.6.2 buildTest.jdl Example	37
4.6.3 Packaging Job Files	40
4.6.4 Deploying Packaged Job Files	40
4.6.5 Running Your Jobs	41
4.6.6 Monitoring Job Results	41
4.6.7 Debugging Jobs	42

5	Policy Elements	45
5.1	Constraints	45
5.2	Facts	45
5.3	Computed Facts	45
6	Using the PlateSpin Orchestrate Client SDK	47
6.1	SDK Requirements	47
6.2	Creating an SDK Client	47
7	Job Architecture	49
7.1	Understanding JDL	49
7.2	JDL Package	50
7.2.1	.sched Files	51
7.3	Job Class	51
7.3.1	Job State Transition Events	51
7.3.2	Handling Custom Events	52
7.4	Job Invocation	53
7.5	Deploying Jobs	53
7.5.1	Using the PlateSpin Orchestrate Development Client	54
7.5.2	Using the zosadmin Command Line Tool	54
7.6	Starting PlateSpin Orchestrate Jobs	55
7.7	Working with Facts and Constraints	55
7.7.1	Grid Objects and Facts	55
7.7.2	Defining Job Elements	56
7.7.3	Job Arguments and Parameter Lists	57
7.8	Using Facts in Job Scripts	58
7.8.1	Fact Values	58
7.8.2	Fact Operations in the Joblet Class	59
7.8.3	Using the Policy Debugger to View Facts	59
7.9	Using Other Grid Objects	59
7.10	Communicating Through Job Events	60
7.10.1	Sending and Receiving Events	60
7.10.2	Synchronization	61
7.11	Executing Local Programs	61
7.11.1	Output Handling	62
7.11.2	Local Users	62
7.11.3	Safety and Failure Handling	63
7.12	Logging and Debugging	63
7.12.1	Creating a Job Memo	63
7.12.2	Tracing	65
7.13	Improving Job and Joblet Robustness	65
7.14	Using an Event Notification in a Job	66
7.14.1	Receiving Event Notifications in a Running Job	66
7.14.2	Event Types	68
8	Job Scheduling	71
8.1	The PlateSpin Orchestrate Job Scheduler Interface	71
8.2	Schedule and Trigger Files	72
8.2.1	Schedule File Examples	72
8.2.2	Trigger File XML Examples	73
8.3	Scheduling with Constraints	75

9	Virtual Machine Job Development	77
9.1	VM Job Best Practices	77
9.1.1	Plan Robust Application Starts and Stops	77
9.1.2	Managing VM Systems	78
9.1.3	Managing VM Images	78
9.1.4	Managing VM Hypervisors	78
9.1.5	VM Job Considerations	78
9.2	Virtual Machine Management	79
9.3	VM Life Cycle Management	80
9.4	Manual Management of a VM Lifecycle	80
9.4.1	Manually Using the zos Command Line	81
9.4.2	Automatically Using the Development Client Job Scheduler	81
9.4.3	Provision Job JDL	81
9.5	Provisioning Virtual Machines	82
9.5.1	Provisioning VMs Using Jobs	84
9.5.2	VM Placement Policy	86
9.5.3	Provisioning Example	87
9.6	Automatically Provisioning a VM	87
9.7	Defining Values for Grid Objects	88
9.7.1	PlateSpin Orchestrate Grid Objects	89
9.7.2	Repository Objects and Facts	90
9.7.3	VmHost Objects and Facts	96
9.7.4	VM Resource Objects and Other Base Resource Facts	101
9.7.5	Physical Resource Objects and Additional Facts	108
10	Complete Job Examples	111
10.1	Accessing Job Examples	111
10.2	Installation and Getting Started	111
10.3	PlateSpin Orchestrate Sample Job Summary	112
10.4	Parallel Computing Examples	113
	demolterator.job	114
	quickie.job	121
10.5	General Purpose Jobs	125
	dgtest.job	126
	failover.job	136
	instclients.job	143
	notepad.job	150
	sweeper.job	155
	whoami.job	162
10.6	Miscellaneous Code-Only Jobs	167
	factJunction.job	168
	jobargs.job	177
A	PlateSpin Orchestrate Client SDK	187
A.1	Constraint Package	187
A.1.1	AndConstraint	187
A.1.2	BetweenConstraint	188
A.1.3	BinaryConstraint	188
A.1.4	Constraint	188
A.1.5	ContainerConstraint	188
A.1.6	ContainsConstraint	188
A.1.7	DefinedConstraint	188
A.1.8	EqConstraint	189
A.1.9	GeConstraint	189

A.1.10	GtConstraint	189
A.1.11	IfConstraint	189
A.1.12	LeConstraint	189
A.1.13	LtConstraint	189
A.1.14	NeConstraint	189
A.1.15	NotConstraint	190
A.1.16	OperatorConstraint	190
A.1.17	OrConstraint	190
A.1.18	TypedConstraint	190
A.1.19	UndefinedConstraint	190
A.1.20	ConstraintException	190
A.2	Datagrid Package	191
A.2.1	GridFile	191
A.2.2	GridFileFilter	191
A.2.3	GridFileNameFilter	191
A.2.4	DGLogger	191
A.2.5	DataGridException	191
A.2.6	DataGridNotAvailableException	192
A.2.7	GridFile.CancelException	192
A.3	Grid Package	192
A.3.1	AgentListener	193
A.3.2	ClientAgent	193
A.3.3	Credential	193
A.3.4	Fact	193
A.3.5	FactSet	194
A.3.6	GridObjectInfo	194
A.3.7	ID	194
A.3.8	JobInfo	194
A.3.9	Message	194
A.3.10	Message.Ack	194
A.3.11	Message.AuthFailure	194
A.3.12	Message.ClientResponseMessage	195
A.3.13	Message.ConnectionID	195
A.3.14	Message.Event	195
A.3.15	Message.GetGridObjects	195
A.3.16	Message.GridObjects	195
A.3.17	Message.JobAccepted	195
A.3.18	Message.JobError	195
A.3.19	Message.JobFinished	196
A.3.20	Message.JobIdEvent	196
A.3.21	Message.JobInfo	196
A.3.22	Message.Jobs	196
A.3.23	Message.JobStarted	196
A.3.24	Message.JobStatus	196
A.3.25	Message.LoginFailed	196
A.3.26	Message.LoginSuccess	197
A.3.27	Message.LogoutAck	197
A.3.28	Message.RunningJobs	197
A.3.29	Message.ServerStatus	197
A.3.30	Node	197
A.3.31	Priority	197
A.3.32	WorkflowInfo	197
A.3.33	ClientOutOfDateException	198
A.3.34	FactException	198
A.3.35	GridAuthenticationException	198
A.3.36	GridAuthorizationException	198
A.3.37	GridConfigurationException	198
A.3.38	GridDeploymentException	198

A.3.39	GridException	198
A.3.40	GridObjectNotFoundException	199
A.4	TLS Package	199
A.4.1	TlsCallbacks	199
A.4.2	PemCertificate	199
A.4.3	TlsConfiguration	199
A.5	Toolkit Package	199
A.5.1	ClientAgentFactory	200
A.5.2	ConstraintFactory	200
A.5.3	CredentialFactory	200

B PlateSpin Orchestrate Job Classes and JDL Syntax 201

B.1	Job Class	201
B.2	Joblet Class	201
B.3	Utility Classes	201
B.4	Built-in JDL Functions and Variables	201
B.4.1	getMatrix()	202
B.4.2	system(cmd)	202
B.4.3	Grid Object TYPE_* Variables	202
B.4.4	The __agent__ Variable	202
B.4.5	The __jobname__ Variable	202
B.4.6	The __mode__ Variable	203
B.5	Job State Field Values	203
B.6	Repository Information String Values	204
B.7	Joblet State Values	204
B.8	Resource Information Values	205
B.9	JDL Class Definitions	205
	AndConstraint()	207
	BinaryConstraint	208
	BuildSpec	209
	CharRange	210
	ComputedFact	211
	ComputedFactContext	212
	Constraint	213
	ContainerConstraint	214
	ContainsConstraint	215
	DataGrid	216
	DefinedConstraint	217
	EqConstraint	218
	Exec	219
	ExecError	220
	FileRange	221
	GeConstraint	222
	GridObjectInfo	223
	GroupInfo	224
	GtConstraint	225
	Job	226
	JobInfo	227
	Joblet	228
	JobletInfo	229
	JobletParameterSpace	230
	LeConstraint	231
	LtConstraint	232
	MatchContext	233
	MatchResult	234
	MatrixInfo	235

MigrateSpec	236
NeConstraint	237
NotConstraint	238
OrConstraint	239
ParameterSpace	240
PolicyInfo	241
ProvisionSpec	242
RepositoryInfo	243
ResourceInfo	244
RunJobSpec	245
ScheduleSpec	246
Timer	247
UndefinedConstraint	248
UserInfo	249
VMHostInfo	250
VmSpec	251

C Understanding Resource Metrics Facts 253

C.1 Resource Facts	253
C.2 Interpreting the Units of Metrics Fact Values	254

D Documentation Updates 257

D.1 July 9, 2009	257
D.2 June 17, 2009 (2.0.2 Release)	257

About This Guide

This *Job Developer Guide and Reference* is a component of the documentation library for PlateSpin® Orchestrator from Novell®. While PlateSpin Orchestrator provides the broad framework and networking tools to manage complex virtual machines and high performance computing resources in a datacenter, this guide explains how to develop grid application jobs and policies that form the basis of PlateSpin Orchestrator functionality. This guide provides developer information to create and run custom PlateSpin Orchestrator jobs. It also helps provide the basis to build, debug, and maintain policies using PlateSpin Orchestrator.

This guide contains the following sections:

- ◆ Chapter 1, “Getting Started With Development,” on page 15
- ◆ Chapter 2, “Advanced Job Development Concepts,” on page 19
- ◆ Chapter 3, “The PlateSpin Orchestrator Datagrid,” on page 25
- ◆ Chapter 4, “Using PlateSpin Orchestrator Jobs,” on page 31
- ◆ Chapter 5, “Policy Elements,” on page 45
- ◆ Chapter 6, “Using the PlateSpin Orchestrator Client SDK,” on page 47
- ◆ Chapter 7, “Job Architecture,” on page 49
- ◆ Chapter 8, “Job Scheduling,” on page 71
- ◆ Chapter 9, “Virtual Machine Job Development,” on page 77
- ◆ Chapter 10, “Complete Job Examples,” on page 111
- ◆ Appendix A, “PlateSpin Orchestrator Client SDK,” on page 187
- ◆ Appendix B, “PlateSpin Orchestrator Job Classes and JDL Syntax,” on page 201
- ◆ Appendix C, “Understanding Resource Metrics Facts,” on page 253
- ◆ Appendix D, “Documentation Updates,” on page 257

Audience

The developer has control of a self-contained development system where he or she creates jobs and policies and tests them in a laboratory environment. When the jobs are tested and proven to function as intended, the developer delivers them to the PlateSpin Orchestrator administrator.

Prerequisite Skills

As data center managers or IT or operations administrators, it is assumed that users of the product have the following background:

- ◆ General understanding of network operating environments and systems architecture.
- ◆ Knowledge of basic Linux* shell commands and text editors.

Documentation Updates

For the most recent version of this *Job Developer Guide and Reference*, visit the [PlateSpin Orchestrator 2.0 Web site \(http://www.novell.com/documentation/pso_orchestrator20/\)](http://www.novell.com/documentation/pso_orchestrator20/).

Additional Product Documentation

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

- ◆ *PlateSpin Orchestrate 2.0 Getting Started Reference*
- ◆ *PlateSpin Orchestrate 2.0 Upgrade Guide*
- ◆ *PlateSpin Orchestrate 2.0 High Availability Configuration Guide*
- ◆ *PlateSpin Orchestrate 2.0 Administrator Reference*
- ◆ *PlateSpin Orchestrate 2.0 VM Client Guide and Reference*
- ◆ *PlateSpin Orchestrate 2.0 Virtual Machine Management Guide*
- ◆ *PlateSpin Orchestrate 2.0 Development Client Reference*
- ◆ *PlateSpin Orchestrate 2.0 Command Line Reference*
- ◆ *PlateSpin Orchestrate 2.0 Server Portal Reference*

Documentation Conventions

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

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

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

Other typographical conventions used in this guide include the following:

Convention	Description
Italics	Indicates variables, new terms and concepts, and book titles. For example, a job is a piece of work that describes how an application can be run in Grid Management on multiple computers.
Boldface	Used for advisory terms such as Note, Tip, Important, Caution, and Warning.
Keycaps	Used to indicate keys on the keyboard that you press to implement an action. If you must press two or more keys simultaneously, keycaps are joined with a hyphen. For example, Ctrl-C. Indicates that you must press two or more keys to implement an action. Simultaneous keystrokes (in which you press down the first key while you type the second character) are joined with a hyphen; for example, press Stop-a. Consecutive keystrokes (in which you press down the first key, then type the second character) are joined with a plus sign; for example, press F4+q.

Convention	Description
Fixed-width	<p>Used to indicate various types of items. These include:</p> <p>Commands that you enter directly, code examples, user type-ins in body text, and options. For example,</p> <pre>cd mydir</pre> <pre>System.out.println("Hello World");</pre> <p>Enter abc123 in the Password box, then click Next.</p> <pre>-keep option</pre> <p>Jobs and policy keywords and identifiers. For example,</p> <pre><run></pre> <pre></run></pre> <p>File and directory names. For example,</p> <pre>/usr/local/bin</pre> <p>Note: UNIX path names are used throughout and are indicated with a forward slash (/). If you are using the Windows platform, substitute backslashes (\) for the forward slashes (/).</p>
Fixed-width italic and <Fixed-width italic>	<p>Indicates variables in commands and code. For example,</p> <pre>zos login <servername> [--user=] [--passwd=] [--port=]</pre> <p>Note: Angle brackets (< >) are used to indicate variables in directory paths and command options.</p>
(pipe)	<p>Used as a separator in menu commands that you select in a graphical user interface (GUI), and to separate choices in a syntax line. For example,</p> <pre>File New</pre> <pre>{a b c}</pre> <pre>[a b c]</pre>
{ } (braces)	<p>Indicates a set of required choices in a syntax line. For example,</p> <pre>{a b c}</pre> <p>means you must choose a, b, or c.</p>
[] (brackets)	<p>Indicates optional items in a syntax line. For example,</p> <pre>[a b c]</pre> <p>means you can choose a, b, c, or nothing.</p>
< > (angle brackets)	<p>Used for XML content elements and tags, and to indicate variables in directory paths and command options. For example,</p> <pre><template></pre> <pre><DIR></pre> <pre>-class <class></pre>

Convention	Description
. . . (horizontal ellipses)	Used to indicate that portions of a code example have been omitted to simplify the discussion, and to indicate that an argument can be repeated several times in a command line. For example, zosadmin [options]optfile.xmlc [...] docfile
plain text	Used for URLs, generic references to objects, and all items that do not require special typography. For example, http://www.novell.com/documentation/index.html The presentation object is in the presentation layer.
ALL CAPS	Used for SQL statements and HTML elements. For example, CREATE statement <INPUT>
lowercase	Used for XML elements. For example, <onevent> Note: XML is case-sensitive. If an existing XML element uses mixed-case or uppercase, it is shown in that case. Otherwise, XML elements are in lowercase.
PlateSpin Orchestrator Server root directory	Where the PlateSpin Orchestrator Server is installed. The PlateSpin Orchestrator executables and libraries are in a directory. This directory is referred to as the PlateSpin Orchestrator Server root directory or <PlateSpin Orchestrator Server_root>.
Paths	UNIX path names are used throughout and are indicated with a forward slash (/). If you are using the Windows platform, substitute backslashes (\) for the forward slashes (/). For example, UNIX: /usr/local/bin Windows: \usr\local\bin
URLs	URLs are indicated in plain text and are generally fully qualified. For example, http://www.novell.com/documentation/index.html
Screen shots	Most screen shots reflect the Microsoft Windows look and feel.

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.

Novell Support

Novell offers a support program designed to assist with technical support and consulting needs. The Novell support team can help with installing and using the Novell product, developing and debugging code, maintaining the deployed applications, providing onsite consulting services, and delivering enterprise-level support.

Getting Started With Development

1

This *Developer Guide* for PlateSpin® Orchestrate from Novell® is intended for individuals acting as PlateSpin Orchestrate job developers. This document discusses the tools and technology required to create discrete programming scripts—called “jobs”—that control nearly every aspect of the PlateSpin Orchestrate product. The guide also explains how to create, debug, and maintain policies that can be associated with jobs running on the PlateSpin Orchestrate Server.

As a job developer, you need your own self-contained, standalone system with full access to your network environment. As a job developer, you might eventually assume all system roles: job creator, job deployer, system administrator, tester, etc. For more information about jobs, see “Jobs” in the *PlateSpin Orchestrate 2.0 Getting Started Reference*.

This section includes the following information:

- ♦ [Section 1.1, “What You Should Know,” on page 15](#)
- ♦ [Section 1.2, “Prerequisites for the Development Environment,” on page 16](#)

1.1 What You Should Know

This section includes the following information:

- ♦ [Section 1.1.1, “Prerequisite Knowledge,” on page 15](#)
- ♦ [Section 1.1.2, “Setting Up Your Development Environment,” on page 16](#)

1.1.1 Prerequisite Knowledge

This guide assumes you have the following background:

- ♦ Sound understanding of networks, operating environments, and system architectures.
- ♦ Familiarity with the Python development language. For more information, see the following online references:
 - ♦ **Python Development Environment (PyDEV):** The [PyDEV plug-in \(http://pydev.sourceforge.net/\)](http://pydev.sourceforge.net/) enables developers to use Eclipse* for Python and Jython development. The plug-in makes Eclipse a more robust Python IDE and comes with tools for code completion, syntax highlighting, syntax analysis, refactoring, debugging, etc.
 - ♦ **Python Reference Manual:** This reference (<http://python.org/doc/2.1/ref/ref.html>) describes the exact syntax and semantics but does not describe the [Python Library Reference, \(http://python.org/doc/2.1/lib/lib.html\)](http://python.org/doc/2.1/lib/lib.html) which is distributed with the language and assists in development.
 - ♦ **Python Tutorial:** This [online tutorial \(http://python.org/doc/2.1/ref/ref.html\)](http://python.org/doc/2.1/ref/ref.html) helps developers get started with Python.
- ♦ Sound understanding of the PlateSpin Orchestrate Job Development Language (JDL).

JDL incorporates compact Python scripts to create job definitions to manage nearly every aspect of the PlateSpin Orchestrate grid. For more information, see [Appendix B, “PlateSpin Orchestrate Job Classes and JDL Syntax,” on page 201](#).
- ♦ Knowledge of basic UNIX shell commands or the Windows command prompt, and text editors.

- ♦ An understanding of parallel computing and how applications are run on PlateSpin Orchestrate infrastructure.
- ♦ Familiarity with on-line PlateSpin Orchestrate API Javadoc as you build custom client applications. For more information see [Appendix A, “PlateSpin Orchestrate Client SDK,” on page 187.](#)
- ♦ Developer must assume both PlateSpin Orchestrate administrative and end-user roles while testing and debugging jobs.

1.1.2 Setting Up Your Development Environment

To set up a development environment for creating, deploying, and testing jobs, we recommend the following procedure:

- 1 Initially set up a simple, easy-to-manage server, agent, and client on a single machine. Even on a single machine, you can simulate multiple servers by starting extra agents (see “[Installing the Orchestrate Agent Only](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*).
- 2 As you get closer to a production environment, your setup might evolve to handle more complex system demands, such as any of the following:
 - ♦ An Orchestrate Server instance deployed on one computer.
 - ♦ An Orchestrate Agent installed on every managed server.
 - ♦ An Orchestrate Development Client installed on your desktop machine.

From your desktop machine, you can build jobs/policies, and then remotely deploy them using zosadmin command line tool. You can then remotely modify the jobs and other grid object through the PlateSpin Orchestrate Development Client.
- 3 Use a version control system, such as Subversion*, to organize and track development changes.
- 4 Put the job version number inside the deployed file. This will help you keep your job versions organized.
- 5 Create make or Ant scripts for bundling and deploying your jobs.

By leveraging the flexibility of the PlateSpin Orchestrate environment, you should not have to write jobs targeted specifically for one hypervisor technology (Xen*, VMware*, etc.).

1.2 Prerequisites for the Development Environment

- ♦ Install the [Java* Development Kit \(https://jdk3d.sun.com/ECom/EComActionServlet.jsessionid=DCA955A842E56492B469230CC680B2E1\)](https://jdk3d.sun.com/ECom/EComActionServlet.jsessionid=DCA955A842E56492B469230CC680B2E1), version 1.5 or later, to create jobs and to compile a Java SDK client in the PlateSpin Orchestrate environment. The PlateSpin Orchestrate installer ships with a Java Runtime Environment (JRE) suitable for running PlateSpin Orchestrate jobs.
- ♦ **Components to write Python-based Job Description Language (JDL) scripts:**
 - ♦ [Eclipse version 3.2.1 or later. \(http://www.eclipse.org/\)](http://www.eclipse.org/).

- ♦ **Development Environment:** Set up your environment according to the guidelines outlined in “[Planning the Orchestrate Server Installation](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*. In general, the installed PlateSpin Orchestrate Server requires 2 (minimum for 100 or fewer managed resources) to 4 gigabytes (recommended for more than 100 managed resources) of RAM.
- ♦ **Network Capabilities:** For Virtual Machine Management, you need a high-speed Gigabit Ethernet. For more information about network requirements, see “[Orchestrate VM Client](#)” and “[VM Hosts](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.
- ♦ **Initial Configuration:** After you install and configure PlateSpin Orchestrate, start in the agent and user auto registration mode as described in “[First Use of Basic PlateSpin Orchestrate Components](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*. As a first-time connection, the server creates an account for you as you set up a self-contained system.

IMPORTANT: Because auto registration mode does not provide high security, make sure you prevent unauthorized access to your network from your work station during development. As you migrate to a production environment, make sure that this mode is deactivated.

Advanced Job Development Concepts

2

This section provides advanced conceptual information to help you create your own PlateSpin® Orchestrate jobs:

- ♦ [Section 2.1, “JDL Job Scripts,” on page 19](#)
- ♦ [Section 2.2, “Understanding TLS Encryption,” on page 21](#)
- ♦ [Section 2.3, “Understanding Job Examples,” on page 21](#)

2.1 JDL Job Scripts

The PlateSpin Orchestrate job definition language (JDL) is an extended and embedded implementation of Python. The PlateSpin Orchestrate system provides additional constructs to control and access the following:

- ♦ Interaction with the infrastructure under management (requesting resources, querying load, etc.)
- ♦ Distributed variable space with job, user and system-wide scoping
- ♦ Extensible event callbacks mechanism
- ♦ Job logging
- ♦ Datagrid for efficient movement of files across the infrastructure.
- ♦ Automatic distribution of parallel operations
- ♦ Failover logic

For more information about the PlateSpin Orchestrate JDL script editor, see [Section 7.2, “JDL Package,” on page 50](#).

The JDL language allows for the scripted construction of logic that can be modified by external parameters and constraints (through one or more associated policies) at the time the job instance is executed. Development of a job with the JDL (Python) language is very straightforward. For a listing of the job, joblet, and utility classes, see [Appendix B, “PlateSpin Orchestrate Job Classes and JDL Syntax,” on page 201](#).

A simple “hello world” Python script example that runs a given number of times (*numTests*) in parallel (subject to resource availability and policy) is shown below:

```
class exampleJob(Job):
    def job_started_event(self):
        print 'Hello world started: got job_started_event'
        # Launch the joblets
        numJoblets = self.getFact("jobargs.numTests")
        pspace = ParameterSpace()
        i = 1
        while i <= numJoblets:
            pspace.appendRow({'name': 'test'+str(i)})
            i += 1
```

```

self.schedule(exampleJoblet, pspace, {})

class exampleJoblet(Joblet):
    def joblet_started_event(self):
        print "Hello from resource%s" % self.getFact("resource.id")

```

This example script contains two sections:

- ♦ The class that extends the job and runs on the server.
- ♦ The class that extends the joblet that will run on any resource employed by this job.

Because the resources are not requested explicitly, they are allocated based on the resource constraints associated with this job. If none are specified, all resources match. The `exampleJoblet` class would typically execute some process or test based on unique parameters.

2.1.1 Principles of Job Operation

Whenever a job is run on the PlateSpin Orchestrate system it undergoes state transition, as illustrated in [Figure 2-1 on page 20](#). In all, there are 11 states. The following four states are important in understanding how constraints are applied on a job's life cycle through policies:

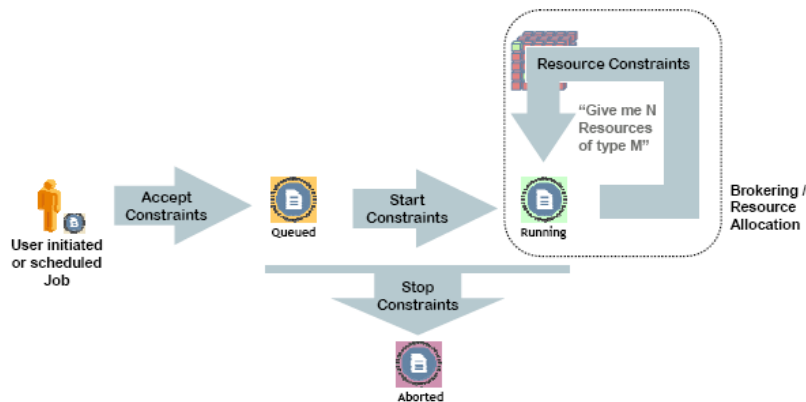
Accept: Used to prevent work from starting; enforces a hard quota on the jobs.

Start: Used to queue up work requests; limits the quantity of jobs or the load on a resource.

Resource: Used to select specific resources.

Stop: Used to abort jobs; provides special timeout or overrun conditions.

Figure 2-1 Constraint-Based Job State Transition



For more information about job life cycle, see [Section 7.3.1, “Job State Transition Events,” on page 51](#).

2.2 Understanding TLS Encryption

Understanding Transport Layer Security (TLS) encryption is particularly important if you reinstall the server and have an old server certificate in either your agent or client user profile similar to `ssh` shared keys. If you have an old certificate, you need to either manually replace it or delete it and allow the client or agent to download the new one from the server using one of the following procedures:

- ♦ **For the Agent:** The TLS certificate is in `<agentdir>/tls/server.pem`. Deleting this certificate will cause the agent, by default, to log a minor warning message and download a new one the next time it tries to connect to the server. This is technically not secure, since the server could be an impersonator. If security is required for this small window of time, then the real server's `<serverdir>/<instancedir>/tls/cert.pem` can be copied to the above `server.pem` file.
- ♦ **For the Client:** The easiest way to update the certificate from the command line tools is to simply answer “yes” both times when prompted about the out-of date certificate. This is, again, not 100% secure, but is suitable for most situations. For absolute security, hand copy the server's `cert.pem` (see above) to `~/.novell/zos/client/tls/<serverIPAddr:Port>.pem`.
- ♦ **For Java SDK clients:** Follow the manual copy technique above to replace the certificate. If the local network is fairly trustworthy, you can also delete the above `~/.novell/.../*.pem` files, which will cause the client to auto-download a new certificate.

2.3 Understanding Job Examples

The following simple examples demonstrate how you can use JDL scripting to manage specific functionality:

- ♦ [Section 2.3.1, “provisionBuildTestResource.job,” on page 21](#)
- ♦ [Section 2.3.2, “Workflow Job Example,” on page 23](#)

To learn about other job examples that are packaged with PlateSpin Orchestrate, see [Chapter 10, “Complete Job Examples,” on page 111](#).

2.3.1 provisionBuildTestResource.job

The following job example illustrates simple scripting to ensure that each of three desired OS platforms might be available in the grid and, if not, it tries to provision them (provided that a VM image matching the OS type exists). The resource Constraint object is created programmatically, so there is no need for external policies.

```
1 class provisionBuildTestResource(Job):
2
3     def job_started_event(self):
4         oslist = ["Windows XP", "Windows 2000", "Windows 2003 Server"]
5         for os in oslist:
6             constraint = EqConstraint()
7             constraint.setFact("resource.os.name")
8             constraint.setValue(os)
9             resources = getMatrix().getGridObjects("resource", constraint)
10            if len(resources) == 0:
11                print "No resources were found to match constraint. \
12 os:%s" % (os)
```

```

13     else:
14         #
15         # Find an offline vm instance or template.
16         #
17         instance = None
18         for resource in resources:
19             if resource.getFact("resource.type") != "Fixed Physical"
and \
20                 resource.getFact("resource.online") == False:
21                 # Found a vm or template. provision it for job.
22                 print "Submitting provisioning request for vm %s." %
(resource)
23                 instance = resource.provision()
24                 print "Provisioning successfully submitted."
25                 break
26         if instance == None:
27             print "No offline vms or templates found for os: %s" %
(os)

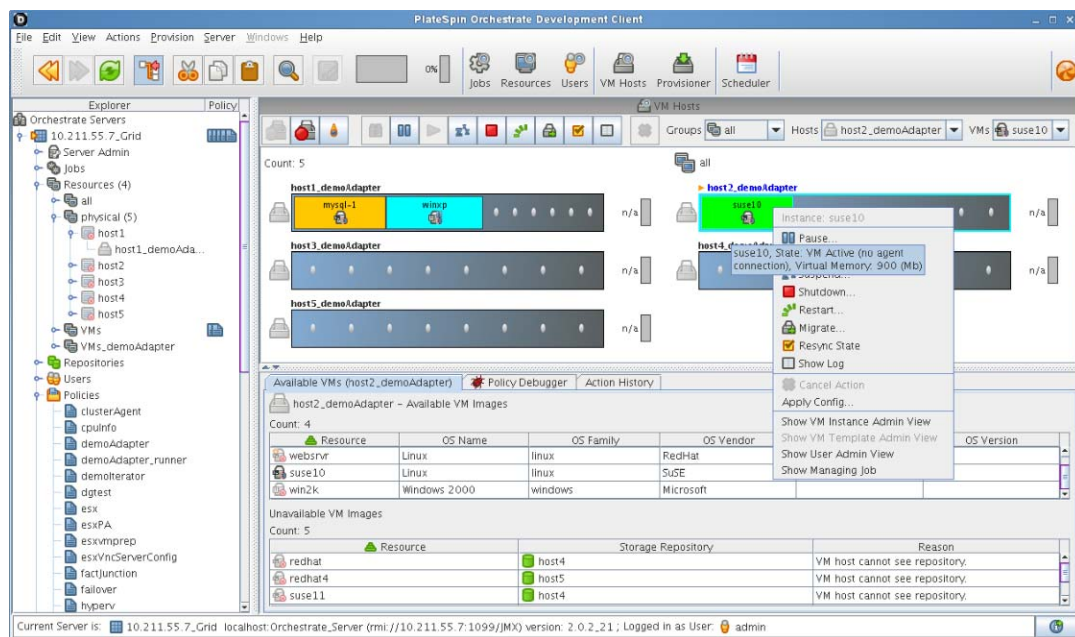
```

It is not necessary to always script resource provisioning. Automatic resource provisioning (“on demand”) is one of the built-in functions of the Orchestrate Server. For example, a job requiring a Windows 2003 Server resource that cannot be satisfied with online resources only needs to have the appropriate facts set in the Orchestrate Development Client; that is, `job.provision.maxcount` is enabled.

This fact could also be set through association with a policy. If it is set up this way, PlateSpin Orchestrate detects that a job is in need of a resource and automatically takes the necessary provisioning steps, including reservation of the provisioned resource.

All provisioned virtual machines and the status of the various hosts are visible in the following view of the Orchestrate Development Client.

Figure 2-2 The PlateSpin Orchestrate Development Client Showing Virtual Machine Management



2.3.2 Workflow Job Example

This brief example illustrates a job that does not require resources but simply acts as a coordinator (workflow) for the buildTest and provision jobs discussed in [Section 4.6, “BuildTest Job Examples,” on page 35](#).

```
1 class Workflow(Job):
2     def job_started_event(self):
3         self.runJob("provisionBuildTestResource", {})
4         self.runJob("buildTest", { "testlist" : "/QA/testlists/production",
5 "buildId": "2006-updateQ1" } )
```

The job starts in line 1 with the `job_started_event`, which initiates [provisionBuildTestResource.job \(page 21\)](#) to ensure all the necessary resources are available, and then starts the [buildTest.jdl Example \(page 37\)](#). This workflow job does not complete until the two subjobs are complete, as defined in lines 3 and 4.

If so desired, this workflow could monitor the progress of subjobs by simply defining new event handler methods (by convention, using the `_event` suffix). The system defines many standard events. Every message received by the job executes the corresponding event handler method and can also contain a payload (a Python dictionary).

The PlateSpin Orchestrate Datagrid

3

This section explains concepts related to the datagrid of the PlateSpin® Orchestrate Server datagrid and specifies many of the objects and facts that are managed in the grid environment:

- ♦ Section 3.1, “Defining the Datagrid,” on page 25
- ♦ Section 3.2, “Datagrid Communications,” on page 27
- ♦ Section 3.3, “datagrid.copy Example,” on page 29

3.1 Defining the Datagrid

Within the PlateSpin Orchestrate environment, the datagrid has three primary functions:

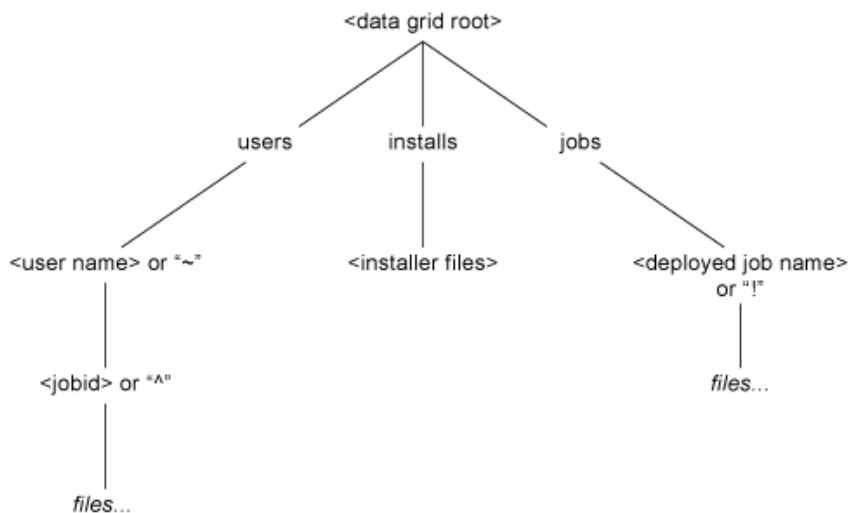
- ♦ Section 3.1.1, “PlateSpin Orchestrate Datagrid Filepaths,” on page 25
- ♦ Section 3.1.2, “Distributing Files,” on page 26
- ♦ Section 3.1.3, “Simultaneous Multicasting to Multiple Receivers,” on page 26

3.1.1 PlateSpin Orchestrate Datagrid Filepaths

The PlateSpin Orchestrate datagrid provides a file naming convention that is used in JDL code and by the PlateSpin Orchestrate CLI for accessing files in the datagrid. The naming convention is in the form of a URL. For more information, see “Jobs” in the *PlateSpin Orchestrate 2.0 Administrator Reference*.

The datagrid defines the root of the namespace as `grid://`, with further divisions under the root as illustrated in the figure below:

Figure 3-1 File Structure of Data Nodes in a Datagrid



The grid URL naming convention is the form `grid://<gridID>/<file path>`. Including the `gridID` is optional and its absence means the host default grid. When writing jobs and configuring a datagrid, you can use the symbol `^` as a shortcut to the `<jobid>` directory either standalone, indicating the current job, or followed by the jobid number to identify a particular job. Likewise, the symbol `!` can be used as a shortcut to the deployed jobs' home directory either standalone, indicating the current jobs' type, or followed by the deployed jobs' name. The symbol `~` is also a shortcut to the user's home directory in the datagrid, either by itself, indicating the current user, or followed by the desired user ID to identify a particular user.

The following examples show address locations in the datagrid using the `zos` command line tool. These examples assume you have logged in using `zos login` to the Orchestrate Server you are using:

- ◆ [“Directory Listing of the Datagrid Root Example” on page 26](#)
- ◆ [“Directory Listing of the Jobs Subdirectory Example” on page 26](#)

Directory Listing of the Datagrid Root Example

```
$ zos dir grid:///
<DIR>      Jun-26-2007 9:42 installs
<DIR>      Jun-26-2007 9:42 jobs
<DIR>      Jun-26-2007 14:26 users
<DIR>      Jun-26-2007 9:42 vms
<DIR>      Jun-26-2007 10:09 warehouse
```

Directory Listing of the Jobs Subdirectory Example

```
$ zos dir grid:///jobs
<DIR>      Jun-26-2007 9:42 cpuInfo
<DIR>      Jun-26-2007 9:42 findApps
<DIR>      Jun-26-2007 9:42 osInfo
<DIR>      Jun-26-2007 9:42 vcenter
<DIR>      Jun-26-2007 9:42 vmHostVncConfig
<DIR>      Jun-26-2007 9:42 vmprep
<DIR>      Jun-26-2007 9:42 vmserver
<DIR>      Jun-26-2007 9:42 vmserverDiscovery
<DIR>      Jun-26-2007 9:42 xen30
<DIR>      Jun-26-2007 9:42 xenDiscovery
<DIR>      Jun-26-2007 9:42 xenVerifier
```

3.1.2 Distributing Files

The PlateSpin Orchestrate datagrid provides a way to distribute files in the absence of a distributed file system. This is an integrated service of PlateSpin Orchestrate that performs system-wide file delivery and management.

3.1.3 Simultaneous Multicasting to Multiple Receivers

The datagrid provides a multicast distribution mechanism that can efficiently distribute large files simultaneously to multiple receivers. This is useful even when a distributed file system is present. For more information, see [Section 3.2, “Datagrid Communications,” on page 27](#).

3.1.4 PlateSpin Orchestrate Datagrid Commands

The following datagrid commands can be used when creating job files. To see where these commands are applied in the PlateSpin Orchestrate Development Client, see “[Typical Use of the Grid](#)”.

Command	Description
cat	Displays the contents of a datagrid file.
copy	Copies files and directories to and from the datagrid.
delete	Deletes files and directories in the datagrid.
dir	Lists files and directories in the datagrid.
head	Displays the first of a datagrid file.
log	Displays the log for the specified job.
mkdir	Makes a new directory in the datagrid.
move	Moves files and directories in the datagrid.
tail	Displays the end of a datagrid file.

3.2 Datagrid Communications

There is no set limit to the number of receivers (nodes) that can participate in the datagrid or in a multicast operation. Indeed, multicast is rarely more efficient when the number of receivers is small. Any type of file or file hierarchy can be distributed via the datagrid.

The datagrid uses both a TCP/IP and IP multicast protocols for file transfer. Unicast transfers (the default) are reliable because of the use of the reliable TCP protocol. Unicast file transfers use the same server/node communication socket that is used for other job coordination datagrid packets are simply wrapped in a generic DataGrid message. Multicast transfers use the persistent socket connection to setup a new multicast port for each transfer.

After the multicast port is opened, data packets are received directly. The socket communication is then used to coordinate packet resends. Typically, a receiver will loose intermittent packets (because of the use of IP multicast, data collisions, etc.). After the file is transferred, all receivers will respond with a bit map of missed packets. The logically ANDing of this mask is used to initiate a resend of commonly missed packets. This process will repeat a few times (with less data to resend on each iteration). Finally, any receiver will still have incomplete data until all the missing pieces are sent in a reliable unicast fashion.

The data transmission for a multicast datagrid transmission is always initiated by the Orchestrate Server. Currently this is the same server that is running the grid.

With the exception of multicast file transfers, all datagrid traffic goes over the existing connection between the agent/client and the server. This is done transparently to the end user or developer. As long as the agent is connected and/or the user is logged in to the grid, the datagrid operations function.

3.2.1 Multicast Example

Multicast transfers are currently only supported through JDL code on the agents. In JDL, after you get the Datagrid object, you can enable and configure multicasting like this:

```
dg.setMulticast(true)
```

Additional multicast tuneables can be set on the object as well, such as the following example:

```
dg.setMulticastRate(20000000)
```

This would set the maximum data rate on the transfer to 20 million bytes/sec. There are a number of other options as well. Refer to the JDL reference for complete information.

The actual multicast copy is initiated when a sufficient number of JDL joblets on different nodes issue the JDL command:

```
dg.copy(...)
```

to actually copy the requested file locally. See the `setMulticastMin` and `setMulticastQuorum` options to change the minimum receiver count and other thresholds for multicasting.

For example, to set up a multicast from a joblet, where the data rate is 30 million bytes/sec, and a minimum of five receivers must request multicast within 30 seconds, but if 30 receivers connect, then start right away, use the following JDL:

```
dg = DataGrid()
dg.setMulticast(true)
dg.setMulticastRate(30000000)
dg.setMulticastMin(5)
dg.setMulticastQuorum(30)
dg.setMulticastWait(30000)
dg.copy('grid:///vms/huge-image.dsk', 'image.dsk')
```

In the above example, if at least five agents running the joblet request the file within the same 30 second period, then a multicast is started to all agents that have requested multicast before the transfer is started. Agents requesting after the cutoff have to wait for the next round. Also, if fewer than 5 agents request the file, then each agent will simply fall back to plain old unicast file copy.

Furthermore, if more than 30 agents connect before 30 seconds is up, then the transfer begins immediately after the 30th request. This is useful for situations where you know how many agents will request the file and want to start as soon as all of them are ready.

3.2.2 Grid Performance Factors

The multicast system performance is dependent on the following factors:

- ♦ **Network Load:** As the load increases, there is more packet loss, which results in more retries.
- ♦ **Number of Nodes:** The more nodes (receivers) there are, the greater the efficiency of the multicast system.
- ♦ **File Size:** The larger the file size, the better. Unless there are a large number of nodes, files less than 2 Mb are probably too small.

- ♦ **Tuning:** The datagrid facility has the ability to throttle network bandwidth. Best performance has been found at about maximum bandwidth divided by 2. Using more bandwidth leads to more collisions. Also the number of simultaneous multicasts can be limited. Finally the minimum receiver size, receiver wait time and quorum receiver size can all be tuned.

Access to the datagrid is typically performed via the CLI tool or JDL code within a job. There is also a Java API in the Client SDK (on which the CLI is implemented). See [“ClientAgent” on page 193](#).

3.2.3 Plan for Datagrid Expansion

When planning your datagrid, you need to consider where you want the Orchestrate Server to store its data. Much of the server data is the contents of the datagrid, including ever-expanding job logs. Every job log can become quite large and quickly exceed its storage constraints.

In addition, every deployed job with its job package—JDL scripts, policy information, and all other associated executables and binary files—is stored in the datagrid. Consequently, if your datagrid is going to grow very large, store it in a directory other than `/opt`.

3.3 datagrid.copy Example

This example fetches the specified source file to the destination. A recursive copy is then attempted if `setRecursive(True)` is set. The default is a single file copy. A multicast also is attempted if `setMulticast(True)` is set. The default is to do a unicast copy. The following example copies a file from the datagrid to a resource, then read the lines of the file:

```
1     datagrid = DataGrid()
2     datagrid.copy("grid:///images/myFile", "myLocalFile")
3     text = open("myLocalFile").readlines()
```

This is an example to recursively copy a directory and its sub directories from the datagrid to a resource:

```
4     datagrid = DataGrid()
5     datagrid.setRecursive(True)
6     datagrid.copy("grid:///testStore/testFiles", "/home/tester/
myLocalFiles")
```

Here’s an example to copy down a file from the job deployment area to a resource and then read the lines of the file:

```
7     datagrid = DataGrid()
8     datagrid.copy("grid:///!myJob/myFile", "myLocalFile")
9     text = open("myLocalFile").readlines()
```

Here are the same examples without using the shortcut characters. This shows the job “myJob” is under the “jobs” directory under the Datagrid root:

```
10    datagrid = DataGrid()
11    datagrid.copy("grid:///jobs/myJob/myFile", "myLocalFile")
12    text = open("myLocalFile").readlines()
```


Using PlateSpin Orchestrate Jobs

4

This section discusses the core job types that can be run by the PlateSpin® Orchestrate Server on grid objects. It also discusses the principles you need to know to run jobs.

- ♦ [Section 4.1, “Resource Discovery,” on page 31](#)
- ♦ [Section 4.2, “Resource Selection,” on page 32](#)
- ♦ [Section 4.3, “Workload Management,” on page 33](#)
- ♦ [Section 4.4, “Policy Management,” on page 34](#)
- ♦ [Section 4.5, “Auditing and Accounting Jobs,” on page 35](#)
- ♦ [Section 4.6, “BuildTest Job Examples,” on page 35](#)

4.1 Resource Discovery

Resource discovery jobs inspect a resource’s environment to set resource facts stored with the resource grid object. These jobs automatically discover the resource attributes (fully extensible facts relating to such things as CPU, memory, storage, bandwidth, load, software inventory) of the resources being managed by the PlateSpin Orchestrate Server. These resource attributes are called “facts.” These facts are used during PlateSpin Orchestrate runtime from within a policy or constraint to select resources that have certain identifiable attributes.

For more information, see [“Walkthrough: Observing Discovery Jobs”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*, and [“Discovering Registered VM Hosts”](#) in the *PlateSpin Orchestrate 2.0 VM Client Guide and Reference*.

4.1.1 Provisioning Jobs

The provisioning jobs included in PlateSpin Orchestrate are used for interacting with VM hosts and repositories for VM life cycle management and for cloning, moving VMs, and other management tasks. These jobs are called “Provisioning Adapters” and are members of the job group called “provisionAdapters.”

With respect to resource discovery, the provisioning jobs are used to discover VM hosts (those resources running a VM technology such as Xen or VmWare) and VM image repositories, as well as VM images residing in those repositories.

For more information, see [Section 9.2, “Virtual Machine Management,” on page 79](#) in this guide.

4.1.2 Resource Discovery Jobs

Some of the commonly used resource discovery jobs include:

- ♦ [“cpuInfo.job” on page 32](#)
- ♦ [“demoInfo.job” on page 32](#)
- ♦ [“findApps.job” on page 32](#)
- ♦ [“osInfo.job” on page 32](#)

cpuInfo.job

Gets CPU information of a resource.

demoInfo.job

Generates the CPU, operating system, and application information for testing.

findApps.job

Finds and reports what applications are installed on the datagrid.

osInfo.job


Gets the operating system of a grid resource. On Linux, it reads the `/proc/cpuinfo`; on Windows, it reads the registry; on UNIX, it executes `uname`.

```
resource.cpu.mhz (integer) e.g., "800" (in Mhz)
resource.cpy.vendor (string) e.g. "GenuineIntel"
resource.cpu.model (string) e.g. "Pentium III"
resource.cpu.family (string) e.g. "i686"
```

4.2 Resource Selection

PlateSpin Orchestrate lets you create jobs that meet the infrastructure scheduling and resource management requirements of your data center, as illustrated in the following figure.

Figure 4-1 *Multi-Dimensional Resource Scheduling Broker*



There are many combinations of constraints and scheduling demands on the system that can be managed by the highly flexible PlateSpin Orchestrate resource broker. As shown in the figure below, many object types are managed by the resource broker. Resource objects are discovered (see [Section 4.1, “Resource Discovery,” on page 31](#)). Other object types such as users and jobs can also be managed. All of these object types have “facts” that define their specific attributes and operational characteristics. PlateSpin Orchestrate compares these facts to requirements set by the administrator for a specific data center task. These comparisons are called “constraints.”

A policy is an XML file that specifies (among other things) constraints and fact values. Policies govern how jobs are dynamically scheduled based on various job constraints. These job constraints are represented in the following figure.

Figure 4-2 Policy-Based Resource Management Relying on Various Constraints

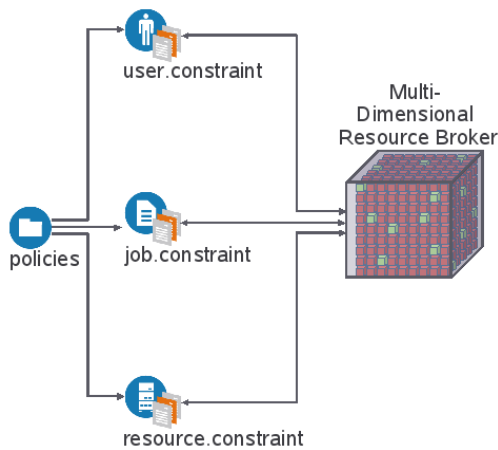
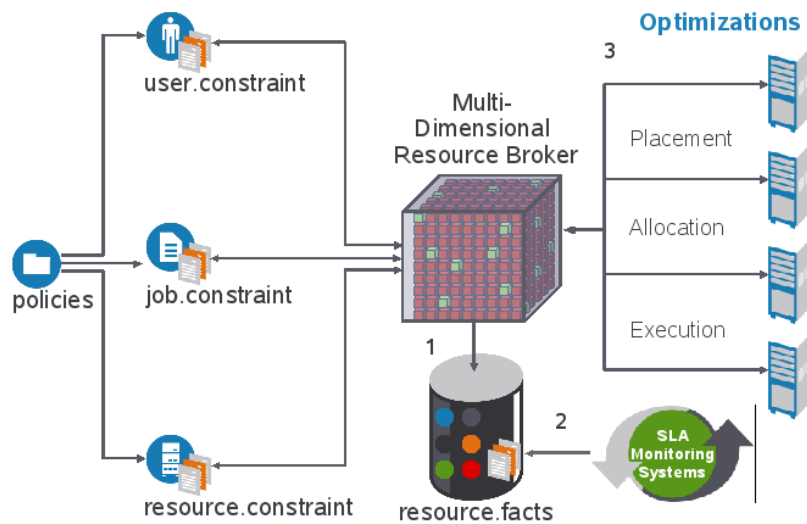


Figure 4-3 Policy-Based Job Management



See [Chapter 8, “Job Scheduling,”](#) on page 71 for examples of scheduling jobs. See also [Section 7.7, “Working with Facts and Constraints,”](#) on page 55.


The Resource Broker function of the Orchestrate Server allocates or “dynamically schedules” resources based on the runtime requirements of a job (for example, the necessary CPU type, OS type, and so on) rather than allocating a specific machine in the data center to run a job. These requirements are defined inside a job policy and are called “resource constraints.” In simpler terms, in order to run a given job, the Resource Broker looks at the resource constraints defined in a job and then allocates an available resource that can satisfy those constraints.

4.3 Workload Management

The Orchestrate Server uses a Provisioning Manager to allocate (assign) and preempt (reassign) resources.

The Provisioning Manager preempts a resource by monitoring the job queue that is waiting for allocation. The manager then compares the job's relative priority to jobs already consuming resources. Higher priority jobs can preempt lower priority jobs.

Figure 4-4 *Workload Management*



Depending on the tasks that applications might require, the Orchestrate Server submits the required jobs to one or more of the connected managed resources to perform specific tasks.

For more information about how job scheduling and provisioning works, see the following sections:

- ♦ [Chapter 8, “Job Scheduling,” on page 71](#)
- ♦ [Section 9.6, “Automatically Provisioning a VM,” on page 87](#)
- ♦ Examples: [dgtest.job \(page 126\)](#) and [factJunction.job \(page 168\)](#).

4.4 Policy Management

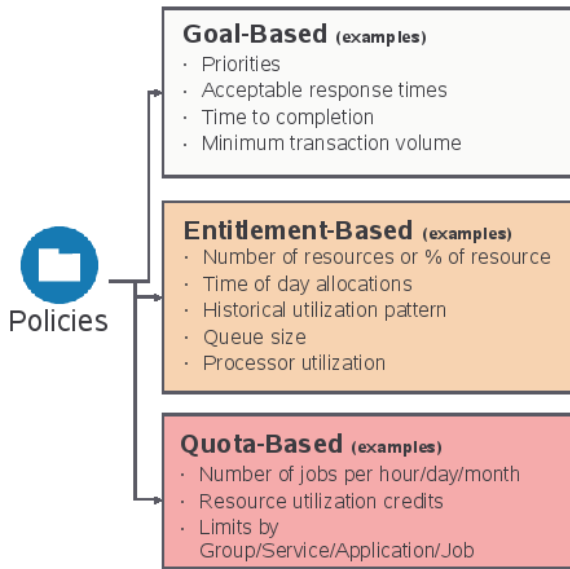
Policies are XML-based files that aggregate the resource facts and constraints that are used to control resources.

Policies are used to enforce quotas, job queuing, resource restrictions, permissions, etc. They can be associated with various grid objects (jobs, users, resources, etc.). The policy example below shows a constraint that limits the number of running jobs to a defined value, while exempting certain users from this limit. Jobs started that exceed the limit are queued until the running jobs count decreases and the constraint passes:

```
<policy>
  <constraint type="start" reason="too busy">
    <or>
      <lt fact="job.instances.active" value="5" />
      <eq fact="user.name" value="canary" />
    </or>
  </constraint>
</policy>
```

Policies can be based on goals, entitlements, quotas, and other factors, all of which are controlled by jobs.

Figure 4-5 Policy Types and Examples



For more information about policies, see [Section 7.7, “Working with Facts and Constraints,”](#) on [page 55](#).

4.5 Auditing and Accounting Jobs

You can create PlateSpin Orchestrate jobs that perform reporting, auditing, and costing functions inside your data center. Your jobs can aggregate cost accounting for assigned resources and perform resource audit trails.

4.6 BuildTest Job Examples

There are many available facts that you can use in creating your jobs. If you find that you need specific kinds of information about a resource or a job, such as the load average of a user or the ID of a job or joblet, chances are that it is already available.

If a fact is not listed, you can create your own facts by creating a `<fact>` element in the job policy. You can also create a fact directly in the JDL job code.

If you want to remember something from one loop to the next or make something available to other objects in the grid, you can set a fact with your own self-defined name.

This section shows an example of a relatively simple working job that performs a set (100) of regression tests on three different platform types. A number of assumptions have been made to simplify this example:

- ♦ Each regression test is atomic and has no dependencies.
- ♦ Every resource is preconfigured to run the tests. Typically, the configuration setup is included as part of the job.

- ♦ The tests are expressed as line entries in a file. PlateSpin Orchestrate has multiple methods to specify parameters. This (/QA/testlists/nightly.dat) is just one example:

```
dir c:/windows
dir c:/windows/system32
dir c:/notexist
dir c:/tmp
dir c:/cygwin
```

To demonstrate the possible functionality for this example, here are some policies that might apply to this example:

- ♦ Only users running tests can use resources owned by their group.
- ♦ To conserve resources, terminate the test after 50 failures.
- ♦ Because the system under test requires a license, prevent more than three of these regression tests from running at one time.
- ♦ To prevent a job backlog, limit the number of queued jobs in the system.
- ♦ To allow the regression test run to tolerate resource failures (for example, unexpected network disconnections, unexpected reboots, and so on), enable automatic failover without affecting the regression run.

The section includes the following information:

- ♦ [Section 4.6.1, “buildTest.policy Example,” on page 36](#)
- ♦ [Section 4.6.2, “buildTest.jdl Example,” on page 37](#)
- ♦ [Section 4.6.3, “Packaging Job Files,” on page 40](#)
- ♦ [Section 4.6.4, “Deploying Packaged Job Files,” on page 40](#)
- ♦ [Section 4.6.5, “Running Your Jobs,” on page 41](#)
- ♦ [Section 4.6.6, “Monitoring Job Results,” on page 41](#)
- ♦ [Section 4.6.7, “Debugging Jobs,” on page 42](#)

4.6.1 buildTest.policy Example

Policies are typically spread over different objects, entities, and groups on the system. However, to simplify the concept, we have combined all policies into this one example that is directly associated with the job.

The arguments available to the job are specified in the in the <jobargs> section (lines 1-11). When the job is run, job arguments are made available as facts to the job instance. The default values of these arguments can be overridden when the job is invoked.

```
1 <policy>
2   <jobargs>
3     <fact name="buildId"
4       type="String"
5       value="02-24-06 1705"
6       description="Build Id to show in memo field" />
7     <fact name="testlist"           type="String"
8       value="/QA/testlists/nightly.dat"
9       description="Path to testlist to use in tests" />
10
11  </jobargs>
```

The `<job>` section (lines 12-25) defines facts that are associated with the job. These facts are used in other policies or by the JDL logic itself. Typically, these facts are aggregated from inherited policies.

```
12 <job>
13   <fact name="max_queue_size"
14     type="Integer"
15     value="10"
16     description="Limit of queued jobs. Any above this limit are not
accepted." />
17   <fact name="max_licenses"
18     type="Integer"
19     value="5"
20     description="License count to limit number of jobs to run
simultaneously. Any above this limit are queued." />
21   <fact name="max_test_failures"
22     type="Integer"
23     value="50"
24     description="To decide to end the job if the number of failures
exceeds a limit" />
25 </job>
```

The `<accept>` (line 26), `<start>` (line 31), and `<continue>` (line 40) constraints control the job life cycle and implement the policy outlined in the example. In addition, allowances are made for “privileged users” (lines 28 and 33) to bypass the accept and start constraints.

```
26 <constraint type="accept" reason="Maximum number of queued jobs has been
reached">
27   <or>
28     <defined fact="user.privileged_user" />
28     <lt fact="job.instances.queued" factvalue="job.max_queue_size" />
29   </or>
30 </constraint>
31 <constraint type="start">
32   <or>
33     <defined fact="user.privileged_user" />
34     <lt fact="job.instances.active" factvalue="job.max_licenses" />
35   </or>
36 </constraint>
```

The `<resource>` constraint (lines 37 and 38) ensures that only resources that are members of the `buildtest` group are used by this job.

```
37 <constraint type="resource">
38   <contains fact="resource.groups" value="buildtest" reason="No
resources are in the buildtest group" />
39 </constraint>
40 <constraint type="continue" >
41   <lt fact="jobinstance.test_failures"
factvalue="job.max_test_failures" reason="Reached test failure limit" />
42 </constraint>
</policy>
```

4.6.2 buildTest.jdl Example

The following example shows how additional resource constraints representing the three test platform types are specified in XML format. These also could have been specified in PlateSpin Orchestrate Development Client.

Setting Resource Constraints

The annotated JDL code represents the job definition, consisting of base Python v2.1 (and libraries) as well as a large number of added PlateSpin Orchestrate operations that allow interaction with the Orchestrate Server:

```
1 import sys,os,time

2 winxp_platform = "<eq fact=\"resource.os.name\" value=\"Windows XP\" />"
3 win2k_platform = "<eq fact=\"resource.os.name\" value=\"Windows 2000\" />"
4 win2003_platform = "<eq fact=\"resource.os.name\" value=\"Windows 2003
Server\" />"
```

Lines 2-4 specify the resource constraints representing the three test platform types (Windows XP, Windows 2000, and Windows 2003) in XML format.

The `job_started_event` in line 6 is the first event delivered to the job on the server. The logic in this method performs some setup and defines the parameter space used to iterate over the tests.

```
5 class BuildTest(Job):

6     def job_started_event(self):
7         self.total_counts = {"failed":0,"passed":0,"run":0}
8         self.setFact("jobinstance.test_failures",0)

9         self.testlist_fn = self.getFact("jobargs.testlist")
10        self.buildId = self.getFact("jobargs.buildId")
11        self.form_memo(self.total_counts)

12        # Form range of tests based on a testlist file
13        filerange = FileRange(self.testlist_fn)
```

Parameter spaces (lines 14-16) can be multidimensional but, in this example, they schedule three units of work (joblets), one for each platform type, each with a parameter space of the range of lines in the (optionally) supplied test file (lines 21, 24 and 27).

```
14        # Form ParameterSpace defining Joblet Splitting
15        pspace = ParameterSpace()
16        pspace.appendDimension("cmd",filerange)

17        # Form JobletSet defining execution on resources
18        jobletset = JobletSet()
19        jobletset.setCount(1)
20        jobletset.setJobletClass(BuildTestJoblet)
```

Within each platform test, a joblet is scheduled for each test line item on each different platform.

```
21        # Launch tests on Windows XP
22        jobletset.setConstraint(winxp_platform)
23        self.schedule(jobletset)

24        # Launch tests on Windows 2000
25        jobletset.setConstraint(win2k_platform)
26        self.schedule(jobletset)

27        # Launch tests on Windows 2003
28        jobletset.setConstraint(win2003_platform)
29        self.schedule(jobletset)
```

The `test_results_event` in line 32 is a message handler that is called whenever the joblets send test results.

```
30 # Event invoked when a Joblet has completed running tests.
31 #
32 def test_results_event(self, params):
33     self.form_memo(params)
```

Creating a Memo Field

In line 37, the `form_memo` method is called to form an informational string to display the running totals for this test. These totals are displayed in the memo field for the job (visible in the Orchestrate Development Client, and Web interface tools). The memo field is accessed through setting the String fact `jobinstance.memo` in line 55.

```
34 #
35 # Update the totals and write totals to memo field.
36 #
37 def form_memo(self, params):
38     # total_counts will be empty at start
39     m = "Build Test BuildId %s " % (self.buildId)
40     i = 0
41     for key in self.total_counts.keys():
42         if params.has_key(key):
43             total = self.total_counts[key]
44             count = params[key]
45             total += count
46             printable_key = str(key).capitalize()
47             if i > 0:
48                 m += ", "
48             else:
49                 if len(m) > 0:
50                     m += ", "
51             m += printable_key + ": %d" % (total)
52             i += 1
53             self.total_counts[key] = total
54     self.setFact("jobinstance.test_failures", self.total_counts["failed"])
55     self.setFact("jobinstance.memo", m)
```

Joblet Definition

As previously discussed, a joblet is the logic that is executed on a remote resource employed by a job, as defined in lines 56-80, below. The `joblet_started_event` in line 60 mirrors the `job_started_event` (line 6) but runs on a different resource than the server.

The portion of the parameter space allocated to this joblet in line 65-66 represents some portion of the total test (parameter) space. The exact breakdown of this is under full control of the administrator/job. Essentially, the size of the “work chunk” in line 67 is a compromise between overhead and retry convenience.

In this example, each element of the parameter space (a test) in line 76 is executed and the exit code is used to determine pass or failure. (The exit code is often insufficient and additional logic must be added to analyze generated files, copy results, or to perform other tasks.) A message is then sent back to the job prior to completion with the result counts.

```

56 #
57 # Define test execution on a resource.
58 #59 class BuildTestJoblet(Joblet):
60     def joblet_started_event(self):
61         passed = 0
62         failed = 0
63         run = 0
64         # Iterate over parameter space assigned to this Joblet
65         pspace = self.getParameterSpace()
66         while pspace.hasNext():
67             chunk = pspace.next()
68             cmd = chunk["cmd"].strip()
69             rslt = self.run_cmd(cmd)
70             print "rslt=%d cmd=%s" % (rslt,cmd)
71             if rslt == 0:
72                 passed +=1
73             else:
74                 failed +=1
75             run += 1
76
self.sendEvent("test_results_event", {"passed":passed,"failed":failed,"run":ru
n})
77     def run_cmd(self,cmd):
78         e = Exec()
79         e.setCommand(cmd)
80         return e.execute()

```

4.6.3 Packaging Job Files

A job package might consist of the following elements:

- ♦ Job description language (JDL) code (the Python-based script containing the bits to control jobs).
- ♦ One or more policy XML files, which apply constraints and other job facts to control jobs. Constraints and facts are discussed [Constraints \(page 45\)](#) and [Facts \(page 45\)](#).
- ♦ Any other associated executables or data files that the job requires. For example, the `cracker.jdl` sample job includes a set of Java code that discovers the user password in every configured agent before the Java class is run.

Or, many discovery jobs that measure performance of Web servers or monitor any other applications can include resource discovery utilities that enable resource discovery.

[Section 7.2, “JDL Package,” on page 50](#) provides more information about job elements.

4.6.4 Deploying Packaged Job Files

After jobs are created, you deploy `.jdl` or multi-element `.job` files to the Orchestrate Server by using any of the following methods:

- ♦ Copying job files into the “hot” Orchestrate Server deployment directory. See [“Deploying a Sample System Job”](#) in the *PlateSpin Orchestrate 2.0 Development Client Reference*.

- Using the Orchestrate Development Client.
- Using the PlateSpin Orchestrate command line (CLI) tools. This process is discussed in “[The zos Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference* and in “[The zosadmin Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*.

4.6.5 Running Your Jobs

After your jobs are deployed, you can execute them by using the following methods:

- **Command Line Interface:** Nearly all PlateSpin Orchestrate functionality, including deploying and running jobs, can be performed using the command line tool, as shown in the following example:

```
zos run buildTest [testlist=mylist]
JobID: paul.buildTest.14
```

More detailed CLI information is available in the zos command line tool.

- **Server Portal:** After PlateSpin Orchestrate is installed, you can use the PlateSpin Orchestrate Server Portal to run jobs. This process is discussed in the *PlateSpin Orchestrate 2.0 Server Portal Reference*.
- **Custom Client:** The PlateSpin Orchestrate toolkit provides an SDK that provides a custom client that can invoke your custom jobs. This process is discussed in [Appendix A, “PlateSpin Orchestrate Client SDK,”](#) on page 187.
- **Built-in Job Scheduler:** The Orchestrate Server uses a built-in Job Scheduler to run deployed jobs. This tool is discussed in [Chapter 8, “Job Scheduling,”](#) on page 71 and in “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*.
- **From Other Jobs:** As part of a job workflow, jobs can be invoked from within other jobs. You integrate these processes within your job scripts as described in the [Chapter 8, “Job Scheduling,”](#) on page 71.

4.6.6 Monitoring Job Results

PlateSpin Orchestrate lets you monitor jobs by using the same methods outlined in [Section 4.6.5, “Running Your Jobs,”](#) on page 41.

Monitoring Jobs from the Command Line

The following example shows the status of the job `ray.buildtest.18` using different monitoring interfaces:

```
zos status -e ray.buildtest.18

Job Status for ray.buildtest.18
-----
                State: Completed          (0 this job)
  Resource Count: 0
Percent Complete: 100%
      Queue Pos: 1 of 1 (initial pos=1)
  Child Job Count: 0                      (0 this job)
```

```

Instance Name: Buildtest
Job Type: buildtest
Memo: Build Test BuildID 02-02-09 1705 , failed: 1, Run: 5,
      Passed: 4
Priority: medium
Arguments: <none>

Submit Time: 02/02/2009 01:46:12
Delayed Start: n/a
Start Time: 02/02/2009 01:46:12
End Time: 01/01/1009 01:46:14
Elapsed Time: 0:00:01
Queue Time: 0:00:00
Pause Time: 0:00:00

Total CPU Time: 0:00:00      (0:00:00 this job)
Total GCycles: 0:00:00      (0:00:00 this job)
Total Cost: $0.0002          ($0.0002 this job)
Burn Rate: $0.0003/hr        (0.0003/hr this job)

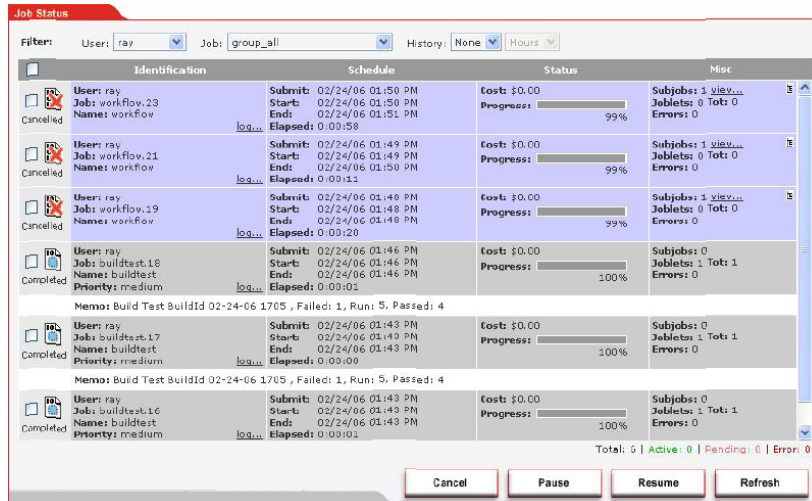
```

The bottom section of the status report shows that you can also monitor job costing metrics, which are quite minimal in this example. More sophisticated job monitoring is possible.

Monitoring Jobs from the Server Portal

You can use the status page of the Server Portal to monitor jobs. In this example, the job memo field is displayed.

Figure 4-6 Server Portal Job Monitoring Example



4.6.7 Debugging Jobs

The following view of the Development Client shows how you can determine that the `buildTest` job was not able to find or match any resources because resources were not added to the `buildtest` group as required by the policy.

Figure 4-7 Debugging Jobs Using the Orchestrate Development Client

The screenshot displays the Orchestrate Development Client interface. At the top, there is a menu bar (File, Edit, View, Actions, Provision, Server, Windows, Help) and a toolbar with icons for Jobs, Resources, Users, VM Hosts, Provisioner, and Scheduler. Below the toolbar, there are filters for Priority (set to 'high'), Username (set to 'All'), and When (set to '15 minutes'). A 'Count: 7' indicator and a 'Toggle Subjob Tree' button are also present.

The main area shows a table of job instances:

Submit Time	Job ID	Instance Name	Start Time	End Time	Elapsed	Σ Resources	Priority	Status	Progress
13:01:45	system.xen30.51	Provision(xen30)	13:01:47	13:03:03	0:01:15	0	low	Completed	
13:03:03	system.findApps...	Scheduler(findApps)	13:03:04	13:03:06	0:00:02	0	high	Completed	
13:03:04	system.cpuinfo.53	Scheduler(cpuinfo)	13:03:04	13:03:05	0:00:01	0	high	Completed	
13:03:04	system.osinfo.54	Scheduler(osinfo)	13:03:04	13:03:06	0:00:02	0	high	Completed	
13:03:04	system.xen30.55	Scheduler(xenDiscovery)	13:03:04	13:03:04	0:00:00	0	medium	Completed	
13:03:04	system.vmserver...	Scheduler(vmserverDisc...	13:03:05	13:03:08	0:00:02	0	medium	Completed	
13:03:05	system.vmBuilder...	Scheduler(vmBuilderDisc...	13:03:05	13:03:08	0:00:03	0	medium	Completed	

Below the job list, the 'Policy Debugger' window is open. It shows a 'Match Context' with 'tszen4_Grid' and 'system.cpuinfo'. The 'Constraint Type' is set to 'resource', and it indicates '1 matching Resources of 1 online'. The 'Constraints' pane shows a tree structure of XML constraints, including a blocking constraint: `<constraint type="resource">`. The 'Referenced Facts' pane shows a table of facts for 'tszen4':

Fact	Type	Value
job.resourcegroup	String	a11
resource.enabled	Boolean	true
resource.groups	String[]	[physical, a11]
resource.id	String	tszen4

The policy debugger shows the blocking constraint, and the tooltip gives the reason. If you drag and drop to add resources to the required group, the job continues quickly with no restart.

Policy Elements

PlateSpin® Orchestrate policy elements are represented in XML. 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:

- ◆ [Section 5.1, “Constraints,” on page 45](#)
- ◆ [Section 5.2, “Facts,” on page 45](#)
- ◆ [Section 5.3, “Computed Facts,” on page 45](#)

5.1 Constraints

The constraint element defines the selection of grid objects such as resources. The required type attribute defines the selection type. Supported types are:

- ◆ resource
- ◆ provision
- ◆ allocation
- ◆ accept
- ◆ start
- ◆ continue
- ◆ vmhost
- ◆ repository

Constraints can also be constructed in JDL and in the Java Client SDK. A JDL constructed constraint can be used for grid search and for scheduling. A Java Client SDK constructed constraint can only be used for grid object search. See also [Section 7.7, “Working with Facts and Constraints,” on page 55](#).

5.2 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.

See the example, `/allTypes.policy`. This example policy has an XML representation for all the fact types.

Facts can also be created and modified in JDL and in the Java Client SDK.

5.3 Computed Facts

Computed facts are used when you want to calculate the value for a fact. Although computed facts are not jobs, they use the same JDL syntax.

To create a new computed fact outside of the Orchestrate Development Client, you can create a file with a `.cfact` extension with the JDL to compute the fact value.

After the new computed fact is created, you deploy it using the same procedures required for jobs (using either the zosadmin command line tool or the Orchestrate Development Client).

The following example shows a computed fact that returns the number of active job instances for a specific job for the current job instance. This fact can be used in an accept or start constraint to limit how many jobs a user can run in the system. The constraint is added to the job policy in which to have the limit. In this example, the start constraint uses this fact to limit the number of active jobs for a user to one:

```
"""
    <constraint type="start" >
      <lt fact="cfact.activejobs"
        value="1"
        reason="You are only allowed to have 1 job running at a time" />
    </constraint>
```

Change JOB_TO_CHECK to define which job is to be limited.

```
"""
JOB_TO_CHECK="quickie"

class activejobs(ComputedFact):

    def compute(self):

        j = self.getContext()
        if j == None:
            # This means computed Fact is executed in a non running
            # job context. e.g., the ZOC fact browser
            print "no job instance"
            return 0
        else:
            # Computed fact is executing in a job context
            user = j.getFact("user.id")
            activejobs = self.getMatrix().getActiveJobs()
            count = 0
            for j in activejobs:
                jobname = j.getFact("job.id")

                # Don't include queued in count !
                state = j.getFact("jobinstance.state.string")
                if jobname == JOB_TO_CHECK \
                    and j.getFact("user.id") == user \
                    and (state == "Running" or state == "Starting"):
                    count+=1

            jobid = j.getFact("jobinstance.id")
            print "jobid=%s count=%d" % (jobid,count)
            return count
```

For another computed fact example, see activejobs.cfact (located in the examples/activejobs.cfact directory).

Using the PlateSpin Orchestrate Client SDK

6

PlateSpin® Orchestrate from Novell® includes a Java* Client SDK in which you can write Java applications to remotely manage jobs. The zos command line tool is written using the Client SDK, as described in [Appendix A, “PlateSpin Orchestrate Client SDK,” on page 187](#). This SDK application can perform the following operations:

- ♦ Login and logout to an Orchestrate Server.
- ♦ Manage the life cycle of a job (run/cancel).
- ♦ Monitor running jobs (get job status).
- ♦ Communicate to a running job using events.
- ♦ Receive events from a running job.
- ♦ Search for grid objects using constraints.
- ♦ Retrieve and modify grid object facts.
- ♦ Datagrid operations (such as copying files to the server and downloading files from the server).

6.1 SDK Requirements

Before you can run the PlateSpin Orchestrate Client SDK, you must perform the following tasks:

1. Install the PlateSpin Orchestrate Client package. For instructions, see [“Installing and Configuring All PlateSpin Orchestrate Components Together”](#) and [“Walkthrough: Launching the PlateSpin Orchestrate Development Client”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.
2. Install JDK 1.4.x or 1.5.x.
3. Examine the two example PlateSpin Orchestrate SDK client applications that are included in the examples directory:
 - ♦ **extranetDemo**: Provides a sophisticated example of launching multiple jobs and listening and sending events to running jobs.
 - ♦ **cracker**: Demonstrates a simple example how to launch a job and listen for events sent from the job to the client application..

6.2 Creating an SDK Client

Use the following procedure to create an SDK client in conjunction with the sample Java code (see [Section A.3.2, “ClientAgent,” on page 193](#)):

- 1 Create ClientAgent instance:

```
// example zos server host is "myserver"

ClientAgent clientAgent = ClientAgentFactory.newClientAgent("myserver");
```

- 2 Use the following user and password example to log in to the Orchestrate Server (see “[Walkthrough: Logging In to the PlateSpin Orchestrate Server](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*):

```
Credential credential =  
CredentialFactory.newPasswordCredential (username,password) ;  
  
clientAgent.login(credential) ;
```

- 3 Run the server operations. In this case, it is the `quickie.jdl` example job (which must have been previously deployed) with no job arguments:

```
String jobID = clientAgent.runJob("quickie",null)
```

- 4 (Optional) Listen for server events using the `AgentListener` interface:

```
clientAgent.addAgentListener(this) ;
```

- 4a Register with the PlateSpin Orchestrate Server to receive job events for the job you started.

```
clientAgent.getMessages (jobID) ;
```

- 5 Log out of the server:

```
clientAgent.logout ()
```


The PlateSpin® Orchestrate Job Scheduler is a sophisticated scheduling engine that maintains high performance network efficiency and quality user service when running jobs on the grid. Such efficiencies are managed through a set of grid component facts that operate in conjunction with job constraints. Facts and constraints operate together like a filter system to maintain both the administrator’s goal of high quality of service and the user’s goal to run fast, inexpensive jobs.

This section explains the following job architectural concepts:

- ◆ [Section 7.1, “Understanding JDL,” on page 49](#)
- ◆ [Section 7.2, “JDL Package,” on page 50](#)
- ◆ [Section 7.3, “Job Class,” on page 51](#)
- ◆ [Section 7.4, “Job Invocation,” on page 53](#)
- ◆ [Section 7.5, “Deploying Jobs,” on page 53](#)
- ◆ [Section 7.6, “Starting PlateSpin Orchestrate Jobs,” on page 55](#)
- ◆ [Section 7.7, “Working with Facts and Constraints,” on page 55](#)
- ◆ [Section 7.8, “Using Facts in Job Scripts,” on page 58](#)
- ◆ [Section 7.9, “Using Other Grid Objects,” on page 59](#)
- ◆ [Section 7.10, “Communicating Through Job Events,” on page 60](#)
- ◆ [Section 7.11, “Executing Local Programs,” on page 61](#)
- ◆ [Section 7.12, “Logging and Debugging,” on page 63](#)
- ◆ [Section 7.13, “Improving Job and Joblet Robustness,” on page 65](#)
- ◆ [Section 7.14, “Using an Event Notification in a Job,” on page 66](#)

7.1 Understanding JDL

The PlateSpin Orchestrate Grid Management system uses an embedded Python-based language for describing jobs (called the Job Definition Language or JDL). This scripting language is used to control the job flow, request resources, handle events and generally interact with the Grid server as jobs proceed.

Jobs run in an environment that expects facts (information) to exist about available resources. These facts are either set up manually through configuration or automatically discovered via discovery jobs. Both the end-user jobs and the discovery jobs have the same structure and language. The only difference is in how they are scheduled.

The job JDL controls the complete life cycle of the job. JDL is a scripting language, so it does not provide compile-time type checking. There are no checks for infinite loops, although various precautions are available to protect against runaway jobs, including job and joblet timeouts, maximum resource consumption, quotas, and limited low-priority JDL thread execution.

As noted, the JDL language is based on the industry standard Python language, which was chosen because of its widespread use for test script writing in QA departments, its performance, its readability of code, and ease to learn.

The Python language has all the familiar looping and conditional operations as well as some powerful operations. There are various books on the language including O'Reilly's *Python in a Nutshell* and *Learning Python*. Online resources are available at <http://www.python.org> (<http://www.python.org>)

Within the Orchestrate Server and grid jobs, JDL not only adds a suite of new commands but also provides an event-oriented programming environment. A job is notified of every state change or activity by calling an appropriately named event handler method.

A job only defines handlers for events it is interested in. In addition to built-in events (such as, `joblet_started_event`, `job_completed_event`, `job_cancelled_event`, and `job_started_event`) it can define handlers for custom events caused by incoming messages. For example, if a job ([Job \(page 226\)](#) class) defines a method as follows:

```
def my_custom_event(self, job, params):
    print "\u201cGot a my_custom event carrying ", params)
```

And the joblet ([Joblet \(page 228\)](#) class) sends an event/message as follows:

```
self.sendEvent("my_custom_event", {"arg1":"one"})
```

NOTE: The event being sent has to be the same name as the defined method receiving the event.

The following line is added to the job log:

```
Got a my_custom event carrying arg1="one"
```

JDL can also define timer events (periodic and one-time) with similar event handlers.

Each event handler can run in a separate thread for parallel execution or can be synchronized to a single thread. A separate thread results in better performance, but also incurs the development expense of ensuring that shared data structures are thread safe.

7.2 JDL Package

The job package consists of the following elements:

- ♦ Job Description Language (JDL) code, consisting of a Python-based script containing the bits to control jobs.
- ♦ An optional policy XML file, which applies constraints and other job facts to control jobs.
- ♦ Any other associated executables or data files that the job requires.

The `cracker.jdl` sample job, for example, includes a set of Java code that discovers the user password in every configured agent before the Java class is run. Or, many discovery jobs, which measure performance of Web servers or monitor any other applications, might include resource discovery utilities that enable resource discovery.

Jobs include all of the code, policy, and data elements necessary to execute specific, predetermined tasks administered either through the PlateSpin Orchestrate Development Client user interface or from the command line. Because each job has specific, predefined elements, jobs can be scripted and delivered to any agent, which ultimately can lead to automating almost any datacenter task.

7.2.1 .sched Files

Job packages also can contain optional XML `.sched` files that describe the scheduling requirements for any job. This file defines when the job is run.

For example, jobs might be run whenever an agent starts up, which is defined in the `.sched` file. The discovery job “[osInfo.job](#)” on [page 32](#) has a schedule XML file that specifies to always run a specified job whenever a specific resource is started and becomes available.

7.3 Job Class

The **Job class** is a representation of a running job instance. This class defines functions for interacting with the server, including handling notification of job state transitions, child job submission, managing joblets and for receiving and sending events from resources and from clients. A job writer defines a subclass of the job class and uses the methods available on the job class for scheduling joblets and event processing.

For more information about the methods this class uses, see [Section 7.3.1, “Job State Transition Events,”](#) on [page 51](#).

The following example demonstrates a job that schedules a single joblet to run on one resource:

```
class Simple(Job):
    def job_started_event(self):
        self.schedule(SimpleJoblet)

class SimpleJoblet(Joblet):
    def joblet_started_event(self):
        print "Hello from Joblet"
```

For the above example, the class `Simple` is instantiated on the server when a job is run either by client tools or by the job scheduler. When a job transitions to the started state, the method `job_started_event` is invoked. Here the `job_started_event` invokes the base class method `schedule()` to create a single joblet and schedule the joblet to run on a resource. The `SimpleJoblet` class is instantiated and run on a resource. A resource is a physical or virtual machine on which the Orchestrator Agent is installed and running and where the Joblet code is executed.

7.3.1 Job State Transition Events

Each job has a set of events that are invoked at the state transitions of a job. On the starting state of a job, the `job_started_event` is always invoked.

The following is a list of job events that are invoked upon job state transitions:

```
job_started_event
job_completed_event
job_cancelled_event
job_failed_event
job_paused_event
job_resumed_event
```

The following is a list of job events that are invoked upon child job state transitions:

```
child_job_started_event
child_job_completed_event
child_job_cancelled_event
child_job_failed_event
```

The following is a list of provisioner events that are invoked upon provisioner state transitions:

```
provisioner_completed_event
provisioner_cancelled_event
provisioner_failed_event
```

The following is a list of joblet events that are invoked as the joblet state transitions:

```
joblet_started_event
joblet_completed_event
joblet_failed_event
joblet_cancelled_event
joblet_retry_event
```

NOTE: Only the `job_started_event` is required; other events are optional.

7.3.2 Handling Custom Events

A job writer can also handle and invoke custom events within a job. Events can come from clients, other jobs, and from joblets.

The following example defines an event handler named `mycustom_event` in a job:

```
class Simple(Job):
    def job_started_event(self):
        ...

    def mycustom_event(self, params):
        dir = params["directory_to_list"]
        self.schedule(MyJoblet, { "dir" : dir } )
```

In this example, the event retrieves a element from the `params` dictionary that is supplied to every custom event. The dictionary is optionally filled by the caller of the event.

The following example invokes the custom event named `mycustom_event` from the PlateSpin Orchestrate client command line tool:

```
zos event <jobid_of_running_job> mycustom_event directory_to_list="/tmp"
```

In this example, a message is sent from the client tool to the job running on the server. The following example invokes the same custom event from a joblet:

```
class SimpleJoblet(Joblet):
    def joblet_started_event(self):
        ...
        self.sendEvent("mycustom_event", {"directory_to_list":"/tmp"})
```

In this example, a message is sent from the joblet running on a resource to the job running on the server. The running job has access to a factset which is the aggregation of the job instance factset (`jobinstance.*`), the deployed job factset (`job.*`, `jobargs.*`), the User factset (`user.*`), the Matrix factset (`matrix.*`) and any jobargs or policy facts supplied at the time the job is started.

Fact values are retrieved using the [GridObjectInfo \(page 223\)](#) functions that the job class inherits.

The following example retrieves the value of the job instance fact `state.string` from the `jobinstance` namespace:

```
class Simple(Job):
    def job_started_event(self):
        jobstate = self.getFact("jobinstance.state.string")
        print "job state=%s" % (jobstate)
```

For further details about each of the events above, see [Section B.1, “Job Class,” on page 201](#).

The following example uses the `joblet_started_event` to determine the resource a Joblet is running on. If you implement the `joblet_started_event` job method, your job is notified when a Joblet has started execution:

```
1 class test(Job):
2     def job_started_event(self):
3         self.schedule(TestJoblet)
4
5     def joblet_started_event(self, jobletNum, resourceId):
6         print "joblet %d is running on %s" % (jobletNum, resourceId)
7
8 class TestJoblet(Joblet):
9     def joblet_started_event(self):
10        import time
11        time.sleep(10)
```

In lines 5 and 6, the `joblet_started_event` is notified when the instance of `TestJoblet` is executing on a resource.

7.4 Job Invocation

Jobs can be started using either the `zos` command line tool, scheduling through a `.sched` file, or manually through the PlateSpin Orchestrate Development Client. Internally, when a job is invoked, an XML file is created. It can be deployed immediately or it can be scheduled for later deployment, depending upon the requirements of the job.

Jobs also can be started within a job. For example, you might have a job that contains JDL code to run a secondary job. Jobs also can be started through the Web portal.

Rather than running jobs immediately, there are many benefits to using the Job Scheduling Manager:

- ♦ Higher priority jobs can be run first and jump ahead in the scheduling priority band.
- ♦ Jobs can be run on the least costly node resources when accelerated performance is not as critical.
- ♦ Jobs can be run on specific types of hardware.
- ♦ User classes can be defined to indicate different priority levels for running jobs.

7.5 Deploying Jobs

A job must be deployed to the Orchestrate Server before it can be run. Deployment to the server is done in either of the following ways:

- ♦ [Section 7.5.1, “Using the PlateSpin Orchestrate Development Client,” on page 54](#)
- ♦ [Section 7.5.2, “Using the `zosadmin` Command Line Tool,” on page 54](#)

7.5.1 Using the PlateSpin Orchestrate Development Client

- 1 In the *Actions* menu, click *Deploy Job*.
- 2 For additional deployment details, see “[Walkthrough: Deploying a Sample Job](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

7.5.2 Using the zosadmin Command Line Tool

From the CLI, you can deploy a component file (.job, .jdl, .sar) or refer to a directory containing job components.

.job files are Java jar archives containing .jdl, .policy, .sched and any other files required by your job. A .sar file is a Java jar archive for containing multiple jobs and policies.

- 1 To deploy a .job file from the command line, enter the following command:

```
>zosadmin deploy <myjob>.job
```
- 2 To deploy a job from a directory where the directory /jobs/myjob contains .jdl, .policy, .sched, and any other files required by your job, enter the following command:

```
>zosadmin deploy /jobs/myjob
```

Deploying from a directory is useful if you want to explode an existing job or .sar file and redeploy the job components without putting the job back together as a .job or .sar file.
- 3 Copy the job file into the “hot” deploy directory by entering the following command:

```
>cp <install dir>/examples/whoami.job <install dir>/deploy
```

As part of an iterative process, you can re-deploy a job from a file or a directory again after specified local changes are made to the job file. You can also undeploy a job out of the system if you are done with it. Use `zosadmin redeploy` and `zosadmin undeploy` to re-deploy and undeploy jobs, respectively.

A typical approach to designing, deploying, and running a job is as follows:

1. Identify and outline the job tasks you want the Orchestrate Server to perform.
2. Use the preconfigured JDL files for specific tasks listed in [Appendix B, “PlateSpin Orchestrate Job Classes and JDL Syntax,”](#) on page 201.
3. To configure jobs, edit the JDL file with an external text editor.
4. Repackage the job as a .jar file.

NOTE: The job could also be packaged and sent as an “exploded” file.

5. Run the zos CLI administration tool to redeploy the packaged job into the Orchestrate Server.
6. Run the job using the zos command line tool.
7. Monitor the results of the job in the PlateSpin Orchestrate Development Client.

Another method to deploy jobs is to edit JDL files through the Orchestrate Development Client. The development client has a text editor that enables you to make changes directly in the JDL file as it is stored on the server ready to deploy. After changes are made and the file is saved using the Orchestrate Development Client, you simply re-run the job without redeploying it. The procedure is useful when you need to fix typos in the JDL file or have minor changes to make in the job functionality.

NOTE: Redeploying a job overwrites any job that has been previously saved on the Orchestrate Server. The Orchestrate Development Client has a *Save File* menu option if you want to preserve JDL modifications you made using the Orchestrate Development Client.

There is no way “undo” a change to a `.jdl` file after the JDL editor in the Orchestrate Development Client has saved the file, nor is there a method for rolling back to a previously deployed version. We recommend that you use an external source code control system such as CVS or SVN for version control.

7.6 Starting PlateSpin Orchestrate Jobs

Jobs can be started by using any of the following options:

- ♦ Running jobs from the `zos` command line (see “[The zos Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*).
- ♦ Running jobs from the PlateSpin Orchestrate Job Scheduler (see “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*).
- ♦ Running jobs from Web applications (see “[Using the PlateSpin Orchestrate Server Portal](#)” in the *PlateSpin Orchestrate 2.0 Server Portal Reference*).
- ♦ Running jobs from within jobs (see “[Using Facts in Job Scripts](#)” on page 58).

7.7 Working with Facts and Constraints

You can incorporate facts and constraints into the custom jobs you create to manage your data center resources using PlateSpin Orchestrate. You should already be familiar with the concepts related to controlling jobs using job facts and constraints. For more information, see the following JDL links:

- ♦ [Job \(page 226\)](#)
- ♦ [Joblet \(page 228\)](#)

This section contains the following topics:

- ♦ [Section 7.7.1, “Grid Objects and Facts,” on page 55](#)
- ♦ [Section 7.7.2, “Defining Job Elements,” on page 56](#)
- ♦ [Section 7.7.3, “Job Arguments and Parameter Lists,” on page 57](#)

7.7.1 Grid Objects and Facts

Every resource and service discovered in an PlateSpin Orchestrate-enabled network is identified and abstracted as an object. Within the PlateSpin Orchestrate management framework, objects are stored within an addressable database called a grid. Every grid object has an associated set of facts and constraints that define the properties and characteristics of either physical or virtual resources. Essentially, by building, deploying, and running jobs on the Orchestrate Server, you can individually change the functionality of any and all system resources by managing an object’s facts and constraints.

The components that have facts include resources, users, jobs, repositories, and vmhosts. The grid server assigns default values to each of the component facts, although they can be changed at anytime by the administrator (unless they are read-only).

However, the developer wants certain constraints to be used for a job and might specify these in the policy. These comprise a set of logical clauses and operators that are compared with the respective component's fact values when the job is run by the Job Scheduling Manager.

Remember, all properties appear in the job context, which is an environment where constraints are evaluated. These constraints provide a multilevel filter for a job in order to ensure the best quality of service the grid can provide.

7.7.2 Defining Job Elements

When you deploy a job, you can include an XML policy file that defines constraints and facts. Because every job is a grid object with its own associated set of facts (`job.id`, etc.), it already has a set of predefined facts, so jobs can also be controlled by changing job arguments at run time.

As a job writer, you define the set of job arguments in the `job args` fact space. Your goal in writing a job is to define the specific elements a job user is permitted to change. These job argument facts are defined in the job XML policy for every given job.

The job argument fact values are passed to a job when the job is run. Consequently, the Orchestrate Server `run` command passes in the job arguments. Similarly, for the job scheduler, you can define which job arguments you want to schedule or run a job. You can also specify job arguments when using the Server Portal.

For example, in the following `quickie.job` example, the number of joblets allowed to run and the amount of sleep time between running joblets are set by the arguments `numJoblets` and `sleeptime` as defined in the policy file for the job. If no job arguments are defined, the client cannot affect the job:

```
...
    # Launch the joblets
    numJoblets = self.getFact("jobargs.numJoblets")
    print 'Launching ', numJoblets, ' joblets'

    self.schedule(quickieJoblet, numJoblets)

class quickieJoblet(Joblet):

    def joblet_started_event(self):
        sleeptime = self.getFact("jobargs.sleeptime")
        time.sleep(sleeptime)
```

To view the complete example, see [quickie.job \(page 121\)](#).

As noted, when running a job, you can pass in a policy file, which is another method the client can use to control job behavior. Policy files can pass in additional constraints to the job, such as how a resource might be selected or how the job runs. The policy file is an XML file defined with the `.policy` extension.

For example, as shown below, you can pass in a policy for the job named `quickie`, with an additional constraint to limit the chosen resources to those with a Linux OS. Suppose a policy file name `linux.policy` in the directory named `/mypolicies` with this content:

```
<constraint type="resource"    <eq fact="resource.os.family" value="linux" /
></constraint>
```

To start the `quickie` job using the additional policy, you would enter the following command:


```
>zos run quickie --policyfile=/mypolicies/linux.policy
```

Creating Constraints Outside of the Constraint Evaluation

When you create constraints, it is sometimes useful to access facts on a Grid object that is not in the context of the constraint evaluation. An example scenario would be to sequence the running of jobs triggered by the Job Scheduler.

In this example, you need to make `job2` run only when all instances of `job1` are complete. To do this, you could add the following start constraint to the `job2` definition:

```
<constraint type="start">
  <eq fact="job[job1].instances.active" value="0"/>
</constraint>
```

Here, the job in the context is `job2`, however the facts on `job1` (`instances.active`) can still be accessed. The general form of the fact name is:

```
<grid_object_type>[<grid_object_name>].rest.of.fact.space
```

PlateSpin Orchestrate supports specific Grid object access for the following grid objects

- Users
- Jobs
- Resources
- VMHosts
- Repositories

Currently, explicit group access is not supported.

7.7.3 Job Arguments and Parameter Lists

Part of a job's static definition might include job arguments. A job argument defines what values can be passed in when a job is invoked. This allows the developer to statically define and control how a job behaves, while the administrator can modify policy values.

You define job arguments in an XML policy file named with the same base name as the job. The example job `cracker.jdl`, for example, has an associated policy file named `cracker.policy`. The `cracker.policy` file contains entries for the `<jobargs>` namespace, as shown in the following partial example from `cracker.policy`.

```
<jobargs>
  <fact name="cryptpw"
    type="String"
    description="Password of abc"
    value="4B3lzcNG/Yx7E"
  />
  <fact name="joblets"
    type="Integer"
    description="joblets to run"
    value="100"
  />
</jobargs>
```

The above policy defines two facts in the `jobargs` namespace for the `cracker` job. One is a `String` fact named `cryptpw` with a default value. The second `jobargs` fact is an integer named `joblets`. Both of these facts have default values so they do not require been set on job invocation. If the default value was omitted, then job would require that the two facts be set on job invocation. The job will not start unless all required job argument facts are supplied at job invocation. The default values of job argument facts can be overridden at job invocation. Job arguments are passed to a job when the job is invoked. This is done in one of the following ways:

- ♦ From the Orchestrate Server `run` command from the CLI, as shown in the following example:

```
>zos run cracker cryptpw="dkslsl"
```

- ♦ From within a JDL job script when invoking a child job, as shown in the following job JDL fragment:

```
self.runjob("cracker", { "cryptpw" : "asdfa" } )
```

- ♦ From the Job Scheduler, either with the Orchestrate Development Client or by a `.sched` file.

7.8 Using Facts in Job Scripts

This section contains the following information:

- ♦ [Section 7.8.1, “Fact Values,” on page 58](#)
- ♦ [Section 7.8.2, “Fact Operations in the Joblet Class,” on page 59](#)
- ♦ [Section 7.8.3, “Using the Policy Debugger to View Facts,” on page 59](#)

7.8.1 Fact Values

Facts can be retrieved, compared against, and written to (if not read-only) from within jobs. Every grid object has a set of accessor and setter JDL functions. For example, to retrieve the `cryptpw` job argument fact in the job example listed in [“Job Arguments and Parameter Lists” on page 57](#), you would write the following JDL code:

```
1 def job_started_event(self):
2     pw = self.getFact("jobargs.cryptpw")
```

In line 2, the function `getFact()` retrieves the value of the job argument fact. `getFact()` is invoked on the job instance `Grid` object.

The following set of JDL grid object functions retrieve facts:

```
getFact()
factExists()
getFactLastModified()
getFactNames()
```

The following set of JDL grid object functions modify fact values (if they are not read-only) and remove facts (if they are not deleteable):

```
setFact
setDateFact
setTimeFact
setArrayFact
setBooleanArrayFact
```

```
setDateArrayFact
setIntegerArrayFact
setTimeArrayFact
setStringArrayFact
deleteFact
```

For more complete information on these fact values, see [GridObjectInfo \(page 223\)](#).

7.8.2 Fact Operations in the Joblet Class

Each Joblet is also a Grid object with its own set of well known facts. These facts are listed in [Section B.2, “Joblet Class,” on page 201](#). An instance of the Joblet class runs on the resource. The Joblet instance on the resource has access to the fact set of the resource where it is running. The resource fact set has no meaning outside of this execution context, because the Joblet can be scheduled to run on any of the resources that match the resource and allocation constraints.

For example, using the `cracker` job example shown in [Section 7.7.3, “Job Arguments and Parameter Lists,” on page 57](#), you would write the following JDL code to retrieve the `cryptpw` job argument fact, the OS family fact for the resource, the Job instanceID fact, and the joblet number:

```
1 class CrackerJoblet(Joblet):
2     def joblet_started_event(self):
3         pw = self.getFact("jobargs.cryptpw")
4         osfamily = self.getFact("resource.os.family")
5         jobid = self.getFact("jobinstance.id")
6         jobletnum = self.getFact("joblet.number")
```

In line 3, the function `getFact()` retrieves the value of the job argument fact. `getFact()` is invoked on the joblet instance grid object. In line 4, the `resource.os.family` fact is retrieved for the resource where the Joblet is being executed. This varies, depending on which resource the Joblet is scheduled to run on. In line 5, the IDfact for the job instance is retrieved. This changes for every job instance. In line 6, the joblet index number for this joblet instance is returned. The index is 0 based.

7.8.3 Using the Policy Debugger to View Facts

The Policy Debugger page of the PlateSpin Orchestrate Development Client provides a table view of all facts in a running or completed job instance. This view includes the Job instance facts (`jobinstance.* namespace`) and the facts from the job context. After you select the Policy Debugger tab in the Job Monitor view, the right side panel displays this fact table. For more information, see “[The Policy Debugger](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*.

7.9 Using Other Grid Objects

Grid objects can be created and retrieved using jobs. This is done when facts from other objects are needed for job decision processing or when joblets are executed on a resource.

The [MatrixInfo \(page 235\)](#) Grid object represents the system and from the `MatrixInfo` object, you can retrieve other grid objects in the system. For example, to retrieve the resource grid object named `webserver` and a fact named `resource.id` from this object, you would enter the following JDL code:

```
webservice = getMatrix().getGridObject(TYPE_RESOURCE, "webservice")
id = webservice.getFact("resource.id")
```

In Line 1, the `ResourceInfo` Grid object for `webservice` is retrieved. The `getMatrix()` built-in function retrieves the `MatrixInfo` object instance. `getGridObject()` is a method on the `MatrixInfo` class. In Line 2, the fact value for the resource fact `resource.id` is retrieved.

The `MatrixInfo` Grid object also provides functions for creating other Grid objects. For more complete information about these functions, see [MatrixInfo \(page 235\)](#).

The `MatrixInfo` object can be used in both `Job` and `Joblet` classes. In the `Joblet` class, `MatrixInfo` cannot create new Grid objects. If your job is required to create Grid objects, you must use `MatrixInfo` in the `Job` class.

7.10 Communicating Through Job Events

JDL events are how the server communicates job state transitions to your job. The required `job_started_event` is always invoked when the job transitions to the starting state.

Likewise, all the other state transitions have JDL equivalents that can be optionally implemented in your job. For example, the `joblet_completed_event` is invoked when a joblet has transitioned to completed. You could implement `joblet_completed_event` to launch another job or joblet or send a custom event to a Client, another job, or another joblet.

You can also use your own custom events for communicating between Client, job, child jobs and joblets.

Every partition of a job (client, job, joblet, child jobs) can communicate directly or indirectly with any other partition of a job by using Events. Events are messages that are communicated to each of the job partitions. For example, a joblet running on a resource can send an event to the job portion running on the server to communicate the completion of a stage of operation.

A job can send an event to a Java Client signalling a stage completion or just to send a log message to display in a client GUI.

Every event carries a dictionary as a payload. You can put any key/values you want to fulfill the requirements of your communication. The dictionary can be empty.

For more information about events are invoked at the state transitions of a job, see [Job \(page 226\)](#) and [Section B.7, “Joblet State Values,” on page 204](#).

7.10.1 Sending and Receiving Events

To send an event from a joblet to a job running on a server, you would input the following:

- 1 The portion in the joblet JDL to send the event:

```
self.sendEvent("myevent", { "message": "hello from joblet" } )
```

- 2 The portion in job JDL to receive the event:

```
def myevent(self, params):
    print "hello from myevent. params=", params
```

To send an event from a job running on the server to a client, you would input the following:

```
self.sendClientEvent("notifyClient", { "log" : "Web server installation
completed" } )
```

In your Java client, you must implement `AgentListener` and check for an Event message.

For testing, you can use the `zos run ... --listen` option to print events from the server. For additional details about the `sendEvent()` and `sendClientEvent()` methods in the [Job \(page 226\)](#) and [Joblet \(page 228\)](#) documentation.

7.10.2 Synchronization

By default, no synchronization occurs on job events. However, synchronization is necessary when you update the same grid objects from multiple events.

In that case, you must put a synchronization wrapper around the critical section you want to protect. The following JDL script is how this is done:

```
1 import synchronize
2 def objects_discovered_event(self, params):
3     print "hello"
4 objects_discovered_event =
synchronize.make_synchronized(objects_discovered_event)
```

Line 1 specifies to use the synchronization wrapper, which requires you to import the `synchronize` package.

Lines 2 and 3 provide the normal definition to an event in your job, while line 4 wraps the function definition with a synchronized wrapper.

7.11 Executing Local Programs

Running local programs is one of the main reasons for scheduling joblets on resources. Although you are not allowed to run local programs on the server side job portion of JDL, there are two ways to run local programs in a joblet:

- 1 Use the built-in `system()` function.

This function is used for simple executions requiring no output or process handling. It simply runs the supplied string as a shell command on the resource and writes `stdout` and `stderr` to the job log.

- 2 Use the `Exec` JDL class.

The `Exec` class provides flexibility in how to invoke executables, to process the output, and to manage the process once running. There is provision for controlling `stdin`, `stdout`, and `stderr` values. `stdout` and `stderr` can be redirected to a file, to the job log, or to a stream object.

`Exec` provides control of how the local program is run. You can choose to run as the agent user or the job user. The default is to run as the job user, but fallback to agent user if the job user does not exist on the resource.

For more information, see [Exec \(page 219\)](#).

7.11.1 Output Handling

The `Exec` (page 219) function provides controls for specifying how to handle `stdout` out `stderr`. By default, `Exec` discards the output.

The following example runs a program that directs `stdout` and `stderr` to the job log:

```
e = Exec()
e.setShellCommand(cmd)
e.writeStdoutToLog()
e.writeStderrToLog()
e.execute()
```

The following example runs a program that directs `stdout` and `stderr` to files and opens the `stdout` file if there is no error in execution:

```
e = Exec()
e.setCommand("ps -aef")
e.setStdoutFile("/tmp/ps.out")
e.setStderrFile("/tmp/ps.err")
result = e.execute()
if result == 0:
    output = open("/tmp/ps.out").read()
    print output
```

7.11.2 Local Users

You can choose to run local programs and have file operations done as the agent user or the job user. The default is to run as the job user, but fallback to agent user if the job user does not exist on the resource. These controls are specified on the job. The `job.joblet.runtype` fact specifies how file and executable operations run in the joblet in behalf of the job user, or not.

The choices for `job.joblet.runtype` are defined in the following table:

Table 7-1 Job Run Type Values

Option	Description
<code>RunAsJobUserFallingBackToNodeUser</code>	Default. This means if the job user exists as a user on the resource, then executable and file operations is done on behalf of that user. By falling back, this means that if the job user does not exist, the agent will still execute the joblet executable and file operation as the agent user. If the executable or file operation still has a permission failure, then the agent user is not allowed to run the local program or do the file operation.
<code>RunOnlyAsJobUser</code>	This means resource can only run the executable or file operation as the job user and will fail immediately if the job user does not exist on the resource. You want to use this mode of operation if you wish to strictly enforce execution and file ownership. You must have your resource setup with NIS or other naming scheme so that your users will exist on the resource.

Option	Description
RunOnlyAsNodeUser	This means the resource will only run executables and do file operations as the agent user.

There is also a fact on the resource grid object that can override the `job.joblet.runttype` fact. The fact `resource.agent.config.exec.asagentuseronly` on the resource grid object can overwrite the `job.joblet.runttype` fact.

This ability to run as the job user is supported by the `enhanced exec` feature of the Orchestrate Agent. A resource might not support PlateSpin Orchestrate enhanced execution of running as job users. If the capability is not supported, the fact `resource.agent.config.exec.enhancedused` is False. This fact is provided so you can create a resource or allocation constraint to exclude such a resource if your grid mixes resource with/without the enhanced exec support and your job requires enhanced exec capabilities.

7.11.3 Safety and Failure Handling

An exception in JDL will fail the job. By default, an exception in the joblet will fail the joblet. The `job.joblet.*` facts provide controls on how many times a failure will fail the joblet. For more information, see [Section 7.13, “Improving Job and Joblet Robustness,” on page 65](#).

```
try :
  < JDL >
except:
  exc_type, exc_value, exc_traceback = sys.exc_info()
  print "Exception:", exc_type, exc_value
```

JDL also provides the `fail()` function on the `Job` and `Joblet` class for failing a job and joblet. The `fail()` function takes an optional reason message.

You would use `fail()` when you detect an error condition and wish to end the job or joblet immediately. Usage of the joblet `fail()` fails the currently running instance of the joblet. The actual failed state of the joblet occurs when the maximum number of retries has been reached.

7.12 Logging and Debugging

The following sections show some examples how jobs can be logged and debugged:

- ♦ [Section 7.12.1, “Creating a Job Memo,” on page 63](#)
- ♦ [Section 7.12.2, “Tracing,” on page 65](#)

7.12.1 Creating a Job Memo

The following job example shows `logExample.jdl` output in the JDL editor of the Orchestrate Development Client.

Figure 7-1 Example Job Displayed in the JDL Editor of the Development Client

```

1 import string, time
2
3
4 # JOB:(facts set in job)
5 # fact name "jobinstance.memo"
6
7 class logExample(Job):
8     def job_started_event(self):
9         numJoblets = 2
10        # Example of text written to job log
11        print "Job Log entry: Running logExample Scheduling %d joblets" % (numJoblets)
12
13        # Set job fact
14        self.setFact("jobinstance.memo", "Running logExample Scheduling " + str(numJoblets) + " joblets")
15
16        # Schedule joblets
17        self.schedule(logExampleJoblet, numJoblets)
18
19
20 # JOBLET: (facts set by joblet instance)
21 # fact name "joblet.memo"
22 # fact name "joblet.instanceName"
23
24 class logExampleJoblet(Joblet):
25     def joblet_started_event(self):
26         sleepTime = 5
27         self.setFact("joblet.memo", "Running logExampleJoblet on " + self.getFact("resource.id"))
28         self.setFact("joblet.instanceName", "myLogExampleJoblet")
29         time.sleep(sleepTime)

```

In the job section of this example (lines 7-17), the fact `jobinstance.memo` (line 14) is set by the job instance. The job log text is emitted on line 11. Both of those are visible in the following example.

Figure 7-2 Example Displaying the `jobinstance.memo` Fact and Job Log Text in the Jobs Monitor View of the Development Client

Job ID	User ID	Instance Name	Start Time	End Time	Elapsed	Priority	Status	Progress
chip.logExample.49	chip	logExample	16:01:15	16:01:20	0:00:05	medium	Completed	<div style="width: 100%;"></div>

Memo
Running logExample Scheduling 2 joblets

```

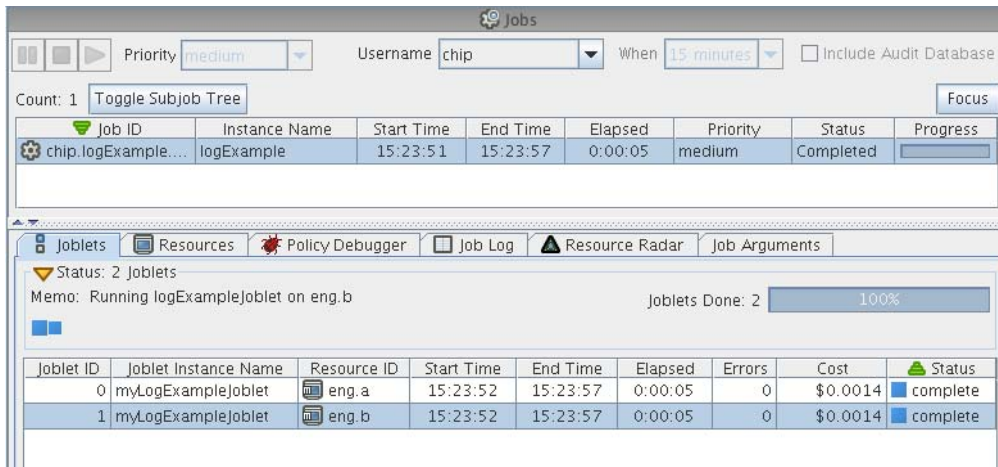
1 Job Log entry: Running logExample Scheduling 2 joblets
2

```

In the `joblet` section of this example (lines 24-29), the fact named `joblet.memo` (line 27) is set by the joblet instance and consists of a brief memo for each joblet. This is typically used for providing detailed explanations, such as the name of the executable being run.

The name of the joblet is specified by the fact named `joblet.instanceName` (line 28). This is typically a simple word displayed in the Development Client joblet column view. The following example shows the joblet facts `joblet.memo` and `joblet.instanceName` in the Development Client.

Figure 7-3 Example of Joblet Facts Displayed in the Development Client



7.12.2 Tracing

There are two facts on the job grid object to turn tracing on or off. The tracing fact writes a message to the job log when a job and/or joblet event is entered and exited. The facts are `job.tracing` and `job.joblet.tracing`. You can turn these on using the Orchestrate Development Client or you can use the `zos run` command tool.

7.13 Improving Job and Joblet Robustness

The job and joblet grid objects provide several facts for controlling the robustness of job and joblet operation.

The default setting of these facts is to fail the job on first error, since failures are typical during the development phase. Depending on your job requirements, you adjust the retry maximum on the fact to enable your joblets either to failover or to retry.

The fact `job.joblet.maxretry` defaults to 0, which means the joblet is not retried. On first failure, the joblet is considered failed. This, in turn, fails the job. However, after you have written and tested your job, you should introduce fault tolerance to the joblet.

For example, suppose you know that your resource application might occasionally timeout due to network or other resource problems. Therefore, you might want to introduce the following behavior by setting facts appropriately:

- ◆ On timeout of 60 seconds, retry the joblet.
- ◆ Retry a maximum of two times. This may cause a retry on another resource matching your resource and allocation constraints.
- ◆ On the third timeout, fail the joblet.

To configure this setup, you use the following facts in either the job policy (using the Orchestrate Development Client to edit the facts directly) or within the job itself:

```
job.joblet.timeout set to 60    job.joblet.maxretry set to 2
```

In addition to timeout, there are different kinds of joblet failures for which you can set the maximum retry. There are forced (job errors) and unforced connection errors. For example, an error condition detected by the JDL code (forced) might require more retries than a network error, which might cause resource disconnections. In the connection failure case, you might want to lower the retry limit because you probably do not want a badly setup resource with connection problems to keep retrying and getting work.

7.14 Using an Event Notification in a Job

Jobs can be notified of a PlateSpin Orchestrate event in two ways.

- ♦ A running Job can subscribe to receive PlateSpin Orchestrate event notifications. (See [Section 7.14.1, “Receiving Event Notifications in a Running Job,” on page 66](#))
- ♦ A Job can be scheduled to start upon an Event notification.

For more information about job scheduling, see [Chapter 8, “Job Scheduling,” on page 71](#) or [“The PlateSpin Orchestrate Job Scheduler” in the *PlateSpin Orchestrate 2.0 Development Client Reference*](#).

7.14.1 Receiving Event Notifications in a Running Job

- ♦ [“Subscribe” on page 66](#)
- ♦ [“Unsubscribe” on page 67](#)
- ♦ [“Callback Method Signature” on page 67](#)
- ♦ [“How an Event Notification Can Start a Job” on page 67](#)

Subscribe

For a job to receive notifications, a job subscribes to an event and must remain running for the notification to occur.

How to subscribe to an event is accomplished using the `subscribeToEvent()` Job method, shown below:

```
def subscribeToEvent( <event Name>, <Job callback method> )
```

In this method, `<event name>` is the string name of the event being subscribed to; `<Job callback method>` is the reference to a Job method. Joblets and globals are not supported.

Example: The following is an example of the `subscribeToEvent()` method.

```
self.subscribeToEvent( "vmhost" ,self.eventHandler)
```

In this example, `vmhost` is the name of the event and `self.eventHandler` is a reference to the callback method.

Unsubscribe

To unsubscribe, use the `unsubscribeFromEvent()` Job method. For the subscribe example shown above, the following is how to unsubscribe.

```
self.unsubscribeFromEvent("vmhost",self.eventHandler)
```

Callback Method Signature

```
def <Job callback method>( self, context):
```

The callback method must be a Job method. The context argument is a dictionary containing name/value pairs. The dictionary contents passed to the callback vary depending on the event type.

Example: The following is an example of the callback method.

```
def eventHandler(self,context):
```

In this method, `context` is the required dictionary argument passed to every callback. The contents of the dictionary vary depending on event type (for details, see [Section 7.14.2, “Event Types,” on page 68](#)).

How an Event Notification Can Start a Job

You can create a schedule using the Job Scheduler or deploy a `.sched` file to start a job on an event notification. For more information, see [Chapter 8, “Job Scheduling,” on page 71](#) or “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*.

The job to be started must match a required job argument signature where the job must define at least one job argument.

The required job argument must be called “context” and be of type Dictionary. The contents of the dictionary vary depending on event type (refer to “Event Types ” below for details).

The contents of `EventResponse.jdl` is an example of a job and policy that can be scheduled on an event notification:

```
1class EventResponse(Job):
2
3    def job_started_event(self):
4        context = self.getFact("jobargs.context")
5
6        print "Context:"
7        keys = context.keys()
8        keys.sort()
9        for k in keys:
10            v = context[k]
11            print "  key(%s) type(%s) value(%s)" % (k,type(v),str(v))
```

Line 4: This line pulls out the job argument for the event context.

Lines 6-11: These lines print out the contents of the context dictionary.

The contents of `EventResponse.policy` are shown below:

```

1<policy>
2  <jobargs>
3    <fact name="context" type="Dictionary"
4      description="Dictionary containing the context for the event " />
5  </jobargs>
6</policy>

```

Lines 3-4: These lines define the required job argument containing the Event context. The running job receives the job argument named `context` with the dictionary completed by the PlateSpin Orchestrate Event Manager with the context that matches the trigger rules.

7.14.2 Event Types

- ♦ “Event Objects” on page 68
- ♦ “Built-in Events” on page 69

Event Objects

Event objects are defined in an XML document and deployed to a server and managed using the Orchestrate Development Client. In the Development Client, these objects are shown in the tree view.

The callback method context argument dictionary contains every grid object type and a value or list of values. The dictionary depends on the event XML definition and the matching grid objects of the `<trigger>` rule.

The following example event file (`vmhost.event`) shows the contents of the dictionary that will be passed as either a jobarg to a job to be scheduled to start, or as a argument to an event callback for a running job.

```

1<event>
2
3  <context>
4    <vmhost />
5    <user>zosSystem</user>
6  </context>
7
8  <trigger>
9    <gt fact="vmhost.resource.loadaverage" value="2" />
10 </trigger>
11
12</event>

```

Lines 3-6: Define the context for the Event object.

Line 4: Defines the match for the trigger rule that iterates over all vmhosts.

Line 5: Defines the context, and contains the user grid object named `zosSystem`.

Assuming that there are 10 vmhosts named "vmhost1, vmhost2, ... vmhost10, but only the first three vmhosts match the trigger rule, the context include a list of the matching vmhosts. In this case, the context dictionary contains the following:

```

1     {
2       "vmhost" : [ "vmhost1", "vmhost2", "vmhost3" ],
3       "user" : "zosSystem",
4       "repository" : "",
5       "resource" : "",
        "job" : ""
    }

```

Line 1: List of the matching VM Hosts that passed the <trigger> rule.

Line 2: The user object that is defined in the <context> XML. In this case, zosSystem.

Lines 3-5: These grid objects are not defined in the context. Their value is empty.

In this example, the dictionary is passed as a job argument to a scheduled job that triggers on the event or is passed to a callback method in a running job that has subscribed to the event.

Built-in Events

Built-in events occur when a managed object comes online or offline or when that object has a health status change. For built-in events, the dictionary contains the name of the grid object type. The value is the name of the grid object.

The PlateSpin Orchestrate built-in events are named as follows:

- ◆ RESOURCE_ONLINE
- ◆ RESOURCE_NEEDS_UPGRADE
- ◆ USER_ONLINE
- ◆ RESOURCE_HEALTH
- ◆ USER_HEALTH
- ◆ VMHOST_HEALTH
- ◆ REPOSITORY_HEALTH

For example, when the resource xen1 comes online, the built-in event called RESOURCE_ONLINE is invoked. Any scheduled jobs are started and any running jobs that have subscribed are invoked. The context dictionary contains the following:

```

{
    "resource" : "xen1"
}

```

The dictionary shown above is passed as a job argument to a scheduled job that triggers on the event or that is passed to a callback method in a running job that has subscribed to the event.

Job Scheduling

8

PlateSpin® Orchestrate from Novell® schedules jobs either start manually using the Job Scheduler or to start programatically using the Job Description Language (JDL). This section contains the following topics:

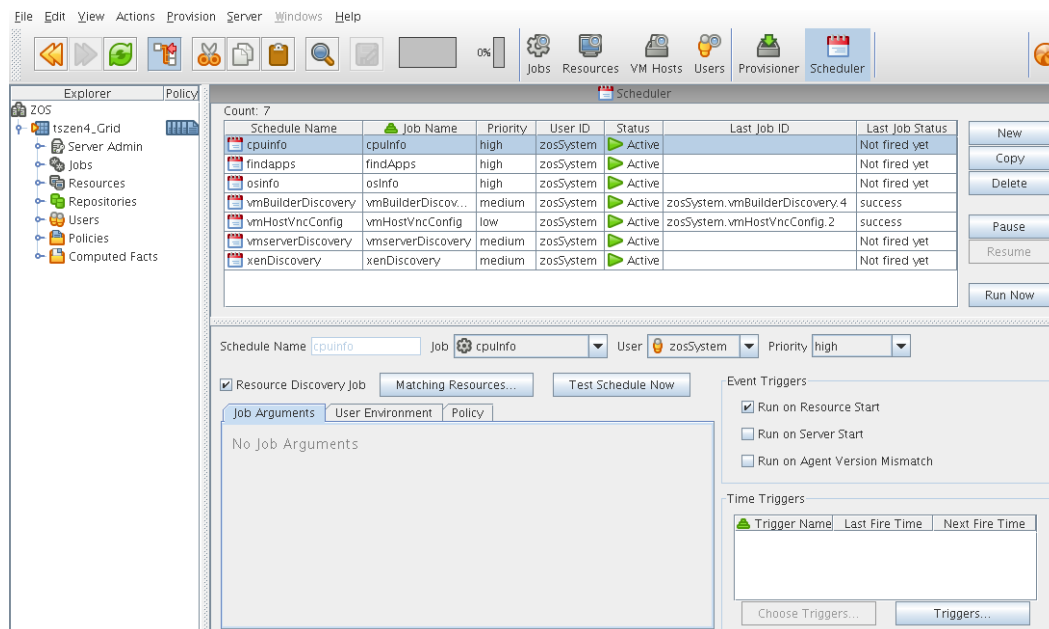
- ◆ [Section 8.1, “The PlateSpin Orchestrate Job Scheduler Interface,” on page 71](#)
- ◆ [Section 8.2, “Schedule and Trigger Files,” on page 72](#)
- ◆ [Section 8.3, “Scheduling with Constraints,” on page 75](#)

8.1 The PlateSpin Orchestrate Job Scheduler Interface

After PlateSpin Orchestrate is enabled with a license, users have access to a built-in job Scheduler. This GUI interface allows jobs to be started periodically based upon user scheduling or when various system or user-defined events occur.

The following figure illustrates the Job Scheduler, with seven jobs staged in the main Scheduler panel.

Figure 8-1 The PlateSpin Orchestrate Scheduler GUI



Jobs are individually submitted and managed using the Job Scheduler as discussed in “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference* and in “[Using the PlateSpin Orchestrate Server Portal](#)” in the *PlateSpin Orchestrate 2.0 Server Portal Reference*.

8.2 Schedule and Trigger Files

In addition to using the Job Scheduler GUI, developers can also write XML files to schedule and trigger jobs to run when triggered by specific events. These files are designated using the `.sched` and `.trig` extensions, and can be included as part of the job archive file (`.job`) or deployed separately.

Everything that you do manually in the Job Scheduler can be automated by creating a `.sched` or `.trig` XML script as part of a job. The XML files enable you to package system and job scheduling logic without using the GUI. This includes setting up cron triggers (for example, running a job at specified intervals) and other triggers that respond to built-in system events, such as resource startup, user startup (that is, login), or user-defined events that trigger on a rule.

For example, the `osInfo` discovery job, which probes a resource for its operating system information, is packaged with a schedule file, as shown in the [“Schedule File Examples” on page 72](#). See also [Section 8.2.2, “Trigger File XML Examples,” on page 73](#).

This section includes the following information:

- ◆ [Section 8.2.1, “Schedule File Examples,” on page 72](#)
- ◆ [Section 8.2.2, “Trigger File XML Examples,” on page 73](#)

8.2.1 Schedule File Examples

A schedule file (`.sched`) can be packaged either within a `.job` archive alongside the `.jdl` file or independently deployed using the `zosadmin` command line utility. Because the XML file defines the job schedule programatically outside of the Orchestrate Development Client, packaging these scripts into jobs is typically a developer task.

This section includes the following information:

- ◆ [“Schedule File Example: osInfo.sched” on page 72](#)
- ◆ [“Schedule File Example: Multiple Triggers” on page 73](#)

Schedule File Example: `osInfo.sched`

The `osinfo.sched` file is packaged with the `osInfo` discovery job, which is deployed as part of the base server. Its purpose is to trigger a run of the `osInfo` job on a resource when the resource comes on line as it logs into the server.

The following shows the syntax of the schedule file that wraps the job:

```
1 <schedule name="osInfo" description="Discover OS info on resources."
active="true">
2   <runjob job="osInfo" user="zosSystem" priority="high" />
3   <triggers>
4     <trigger name="RESOURCE_ONLINE" />
5   </triggers>
6 </schedule>
```

Line 1: Defines a new schedule named `osinfo`, which is used to schedule a run of the job `osInfo`. If the job (in this case, `osinfo`) is not deployed, the deployment returns a “Job is not deployed” error.

Line 2: Instructs the schedule to run the named job (`osinfo`) by the named user (`zosSystem`) using a defined priority (`high`). If the user (in this case, `zosSystem`) does not exist, the deployment returns a “...” error.

NOTE: Only the PlateSpin Orchestrate users belonging to the `administrator` group can assign priorities higher than `medium`. Assigning a higher priority than specified by the `user.priority.max` fact defaults to a priority equal to `user.priority.max` when the job runs.

Line 3-5: Defines the triggers (in this case only one trigger, an event, `RESOURCE_ONLINE`) that initiates the job.

Schedule File Example: Multiple Triggers

A schedule can include one or more triggers. The following example shows the syntax of a schedule file that has two cron triggers for scheduling a job:

```
1 <schedule name="ReportTwice" active="true">
2   <runjob job="jobargs" user="JohnD" priority="medium" />
3   <triggers>
4     <trigger name="DailyReportTrigger" />
5     <trigger name="NightlyReportTrigger" />
6   </triggers>
7 </schedule>
```

Line 1: Defines the schedule name, deployed condition, and description (if any).

Line 2: Instructs the schedule to run the named job (`jobargs`) by the named user (`JohnD`) using a defined priority (`medium`).

Lines 3-6: Defines the triggers (in this case two occurring time triggers) that initiate the job.

8.2.2 Trigger File XML Examples

Trigger files define when a job and how often schedule fires. This can happen when a defined event occurs, when a defined amount of time passes, or when a given point in time is reached. so that one or more triggers can be associated with a job schedule (`.sched`). You can create these triggers yourself in XML format and deploy them, or you can edit and choose them in the Job Scheduler, which automatically deploys them. This section includes examples to show you the syntax of different trigger files.

- ◆ [“XML Example: Event Trigger” on page 73](#)
- ◆ [“XML Example: Interval Time Trigger” on page 74](#)
- ◆ [“XML Example: Cron Expression Trigger” on page 74](#)

XML Example: Event Trigger

An event trigger starts a job when a defined event occurs. Several built-in event triggers (such as events that occur when a managed object comes online or offline or has a health status change) are available in the Trigger chooser of the Job Scheduler along with any user-defined Events:

- ◆ `RESOURCE_ONLINE`
- ◆ `RESOURCE_OFFLINE`
- ◆ `USER_ONLINE`

- ◆ USER_OFFLINE
- ◆ RESOURCE_HEALTH
- ◆ USER_HEALTH
- ◆ VMHOST_HEALTH
- ◆ REPOSITORY_HEALTH

When deployed as triggers in a schedule, built-in events do not generate a `.trig` file, as other triggers do.

You can also associate an Event object with a job schedule (see “Event Triggers” in “The PlateSpin Orchestrate Job Scheduler” in the *PlateSpin Orchestrate 2.0 Development Client Reference*). You can define an Event object in an XML document, deploy it to a server, and then manage it with the PlateSpin Orchestrate Development Client.

The following example, `PowerOutage.trig` shows the XML format for a trigger that references an event object.

```
1 <trigger name="PowerOutage" description="Fires when UPS starts at power
outage">
2   <event value="UPS_interrupt"/>
3 </trigger>
```

Line 1: Defines the trigger name and description.

Line 2: Defines the event object chosen for the trigger.

For more information about events, see [Section 7.14, “Using an Event Notification in a Job,”](#) on page 66.

XML Example: Interval Time Trigger

The following example, `EveryMin1Hr.trig` shows the XML format for a trigger that uses the system clock to define (in seconds) how soon the schedule is to start, how often the schedule is to repeat, and how many times the schedule is to be repeated:

```
1 <trigger name="EveryMin1Hr" description="Fires every minute for one hour">
2   <interval startin="600" interval="60" repeat="60"/>
3 </trigger>
```

Line 1: Defines the trigger name and description.

Lines 2-3: Defines how soon the schedule is to start, how often the schedule is to repeat, and how many times the schedule is to be repeated

XML Example: Cron Expression Trigger

The following example, `NoonDaily.trig` shows the XML format for a trigger that uses a Quartz Cron expression to precisely define when an event is to fire.

```
1 <trigger name="NoonDaily" description="Fires every day at noon">
2   <cron value="0 0 12 * * ?"/>
3 </trigger>
```

Line 1: Defines the trigger name and description.

Lines 2-3: Defines the cron expression to be used by the schedule. Cron expressions are used to precisely define the future point in time when the schedule is to fire. For more information, see “[Understanding Cron Syntax in the Job Scheduler](#)” in “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*.

8.3 Scheduling with Constraints

The constraint specification of the policies is comprised of a set of logical clauses and operators that compare property names and values. The grid server defines most of these properties, but they can also be arbitrarily extended by the user/developer.

All properties appear in the job context, which is an environment where constraints are evaluated. Compound clauses can be created by logical concatenation of earlier clauses. A rich set of constraints can thus be written in the policies to describe the needs of a particular job. However, this is only part of the picture.

Constraints can also be set by an administrator via deployed policies, and additional constraints can be specified by jobs to further restrict a particular job instance. The figure below shows the complete process employed by the Orchestrate Server to constrain and schedule jobs.

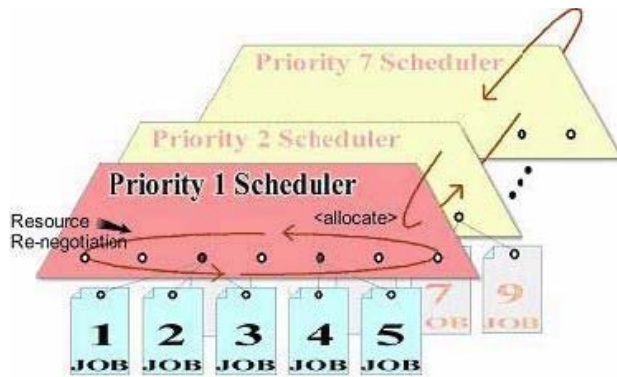
When a user issues a work request, the user facts (user.* facts) and job facts (job.* facts) are added to the job context. The server also makes all available resource facts (resource.* facts) visible by reference. This set of properties creates an environment in which constraints can be executed. The scheduler applies a logic ANDing of job constraints (specified in the policies), grid policy constraints (set on the server), optionally additional user defined constraints specified on job submission, and optional constraints specified by the resources.

This procedure results in a list of matching resources. The PlateSpin Orchestrate solution returns three lists:

- ◆ Available resources
- ◆ Pre-emptable resources (nodes running lower priority jobs that could be suspended)
- ◆ Resources that could be “stolen” (nodes running lower-priority jobs that could be killed)

These lists are then passed to the resource allocation logic where, given the possible resources, the ordered list of desired resources is returned together with information on the minimum acceptable allocation. The scheduler uses this information to appropriate resources for all jobs within the same priority group. Because the scheduler is continually re-evaluating the allocation of resources, the job policies forms part of the scheduler's real-time algorithm, thus providing an extremely versatile and powerful scheduling mechanism.

Figure 8-2 Job Scheduling Priority



Although job scheduling might appear complex, it is very easy to use for an end user. For example, a job developer might write just a few lines of policy code to describe a job to require a node with a x86 machine, greater than 512 MB of memory, and a resource allocation strategy of minimizing execution time. Below is an example.

```
<constraint type="resource">
  <and>
    <eq fact="cpu.architecture" value="x86" />
    <gt fact="memory.physical.total" value="512" />
  </and>
</constraint>
```

Virtual Machine Job Development

9

This section explains the following concepts related to developing virtual machine (VM) management jobs with PlateSpin® Orchestrate:

- ◆ [Section 9.1, “VM Job Best Practices,” on page 77](#)
- ◆ [Section 9.2, “Virtual Machine Management,” on page 79](#)
- ◆ [Section 9.3, “VM Life Cycle Management,” on page 80](#)
- ◆ [Section 9.4, “Manual Management of a VM Lifecycle,” on page 80](#)
- ◆ [Section 9.5, “Provisioning Virtual Machines,” on page 82](#)
- ◆ [Section 9.6, “Automatically Provisioning a VM,” on page 87](#)
- ◆ [Section 9.7, “Defining Values for Grid Objects,” on page 88](#)

9.1 VM Job Best Practices

This section discusses some of VM job architecture best practices to help you understand and get started developing VM jobs:

- ◆ [Section 9.1.1, “Plan Robust Application Starts and Stops,” on page 77](#)
- ◆ [Section 9.1.2, “Managing VM Systems,” on page 78](#)
- ◆ [Section 9.1.3, “Managing VM Images,” on page 78](#)
- ◆ [Section 9.1.4, “Managing VM Hypervisors,” on page 78](#)
- ◆ [Section 9.1.5, “VM Job Considerations,” on page 78](#)

9.1.1 Plan Robust Application Starts and Stops

An application is required for a service, and a VM is provisioned on its behalf. As part of the provisioning process, the VM’s OS typically must be prepared for specific work; for example, NAS mounts, configuration, and other tasks. The application might also need customizing, such as configuring file transfer profiles, client/server relationships, and other tasks.

Then, the application is started and its “identity” (IP address, instance name, and other identifying characteristics) might need to be transferred to other application instances in the service, or a load balancer).

If the Orchestrate Server loses the job/joblet communication state machine, such as when a server failover or job timeout occurs, all of the state information must be able to be recovered from “facts” that are associated with the server. This kind of job should also work in a disaster recovery mode, so it should be implemented in jobs regularly when relevant services from Data Center A must be started in Data Center B in a DR case. These jobs require special precautions.

9.1.2 Managing VM Systems

A series of VMs must typically be provisioned in order to run system-wide maintenance tasks. Because there might not be enough resources to bring up every VM simultaneously, you might consider running discovery jobs to limit how many resources (RAM, cores, etc.) that can be used at any given time. Then, you should consider running a task that writes a consolidated audit trail.

9.1.3 Managing VM Images

Similar to how the job `installagent` searches for virtual machine grid objects using specified Constraints and runs a VM operation (`installAgent`) on the VMs that are returned, a PlateSpin Orchestrate image must be modified when the VM is not running. Preferably, this should occur without having to provision the VM itself.

9.1.4 Managing VM Hypervisors

The management engine (“hypervisor”) underlying the host server must be “managed” while a VM is running. For example, VM memory or CPU parameters must be adjusted on behalf of a monitoring job or a Development Client action.

9.1.5 VM Job Considerations


In some instances, some managed resources might host VMs that do not contain an Orchestrate Agent. Such VMs can only be controlled by administrators interacting directly with them.

Long-running VMs can be modified or migrated while the job managing the VM is not actively interacting with it. If you have one joblet running on the container and one inside the VM, that relationship might have to be re-established.

9.2 Virtual Machine Management

The PlateSpin Orchestrate provisioning manager provides the ability to manage the use of virtual machines, as shown in the following figure:

Figure 9-1 VM Management



For more information about managing virtual machines, see the *PlateSpin Orchestrate 2.0 VM Client Guide and Reference*.

While PlateSpin Orchestrate enables you manage many aspects of your virtual environment, as a developer, you can create custom jobs that do the following tasks:

- ♦ **Create and clone VMs:** These jobs Create virtual machine images to be stored or deployed. They also create templates for building images to be stored or deployed (see “VM Instance:” on page 83 and “VM Template:” on page 83).
- ♦ **Discover resources that can be used as VM hosts.**
- ♦ **Provision, migrate, and move VMs:** Virtual machine images can be moved from one physical machine to another.
- ♦ **Provide checkpoints, restoration, and re-synchronization of VMs:** Snapshots of the virtual machine image can be taken and used to restore the environment if needed. For more information, refer to the documentation for your hypervisor or contact technical support organization for that hypervisor.
- ♦ **Monitor VM operations:** Jobs can start, shut down, suspend and restart VMs.
- ♦ **Manage on, off, suspend, and restart operations.**

9.3 VM Life Cycle Management

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

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

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

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

9.4 Manual Management of a VM Lifecycle

The example provided in this section is a general purpose job that only provisions a resource.

You might use a job like this, for example, each day at 5:00 p.m. when your accounting department requires extra SAP servers to be available. As a developer, you would create a job that provisions the required VMs, then use the PlateSpin Orchestrate Scheduler to schedule the job to run every day at the time specified.

In this example, the `provision` job retrieves the members of a resource group (which are VMs) and invokes the `provision` action on the VM objects. For an example of a provision job JDL, see [Section 9.4.3, “Provision Job JDL,” on page 81](#).

To setup to create the `provision.job`, use the following procedure:

- 1 Create your VMs and follow the discovery process in the Development Client so that the VMs are contained in the PlateSpin Orchestrate inventory.

- 2 In the Development Client, create a Resource Group called `sap` and add the required VMs as members of the group.
- 3 Given the `.jdl` and `.policy` below you would create a `.job` file (jar them):


```
>jar cvf provision.job provision.jdl provision.policy
```
- 4 Deploy the `provision.job` file to the Orchestrate Server using either the Development Client or the `zosadmin` command line.

To run the job, use either of the following procedures:

- ♦ [Section 9.4.1, “Manually Using the zos Command Line,” on page 81](#)
- ♦ [Section 9.4.2, “Automatically Using the Development Client Job Scheduler,” on page 81](#)

9.4.1 Manually Using the zos Command Line

- 1 At the command line, enter:

```
>zos login <zos server>
>zos run provision VmGroup="sap"
```

For more complete details about entering CLI commands, see “[The zos Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*.

9.4.2 Automatically Using the Development Client Job Scheduler

- 1 In the Development Client, create a *New* schedule.
- 2 Fill in the job name (`provision`), user, priority.
- 3 For the jobarg `VmGroup`, enter `sap`.
- 4 Create a Trigger for the time you want this job to run.
- 5 Save the Schedule and enable it by clicking *Resume*.

You can manually force scheduling by clicking *Test Schedule Now*.

For more complete details about using the Job Scheduler, see “[The PlateSpin Orchestrate Job Scheduler](#)” in the *PlateSpin Orchestrate 2.0 Development Client Reference*. You can also refer to [Section 9.6, “Automatically Provisioning a VM,” on page 87](#) in this guide.

9.4.3 Provision Job JDL

```
"""Job that retrieves the members of a supplied resource group and invokes the
provision action on all members. For more details about this class, see Job
(page 226). See also ProvisionSpec (page 242).
```

The members must be VMs.

```
"""
```

```
class provision(Job):
    def job_started_event(self):
        # Retrieves the value of a job argument supplied in
```

```

# the 'zos run' or scheduled run.
VmGroup = self.getFact("jobargs.VmGroup")

#
# Retrieves the resource group grid object of the supplied name.
# The job Fails if the group name does not exist.
#
group = getMatrix().getGroup(TYPE_RESOURCE, VmGroup)
if group == None:
    self.fail("No such group '%s'." % (VmGroup))

#
# Gets a list of group members and invokes a provision action on each
one.
#
members = group.getMembers()
for vm in members:
    vm.provision()
    print "Provision action requested for VM '%s'" %
(vm.getFact("resource.id"))

Job Policy:
<!--
    The policy definition for the provision example job.

    This specifies the job argument VmGroup' which is required
-->
<policy>

    <jobargs>

        <fact name="VmGroup"
            type="String"
            description="Name of a VM resource group whose members will be
provisioned"
            />

    </jobargs>

</policy>

```

9.5 Provisioning Virtual Machines

VM provisioning adapters run just like regular jobs on PlateSpin Orchestrate. The system can detect a local store on each VM host and if a local disk might contain VM images. The provisioner puts in a request for a VM host. However, before a VM is brought to life, the system pre-reserves that VM for exclusive use.

That reservation prevents a VM from being stolen by any other job that's waiting for a resource that might match this particular VM. The constraints specified to find a suitable host evaluates machine architectures, CPU, bit width, available virtual memory, or other administrator configured constraints, such as the number of virtual machine slots.

This process provides heterogeneous virtual machine management using the following virtual machine adapters (also called “provisioning adapters”):

- ♦ **Xen Adapter:** For more information, see [XenSource*](http://www.xensource.com/) (<http://www.xensource.com/>).
- ♦ **Virtual Center 1.3.x or later, and Virtual Center 2.0.1 or later:** For more information, see [VMware](http://www.vmware.com) (<http://www.vmware.com>).
- ♦ **Hyper-V:** For more information, see [Microsoft*](http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx) [Windows Server 2008 Virtualization with Hyper-V](http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx) (<http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx>).
- ♦ **VMware Server 1.x:** For more information, see [VMware](http://www.vmware.com) (<http://www.vmware.com>).
- ♦ **ESX 3.0.x and 3.5.x:** For more information, see [VMware](http://www.vmware.com) (<http://www.vmware.com>).

NOTE: Novell considers the PlateSpin Orchestrate provisioning technology used in conjunction with VMware ESX or VMware Server as experimental because it has not yet been fully tested.

For more information, see *PlateSpin Orchestrate 2.0 Virtual Machine Management Guide* “Provisioning a Virtual Machine” in the *PlateSpin Orchestrate 2.0 Virtual Machine Management Guide*.

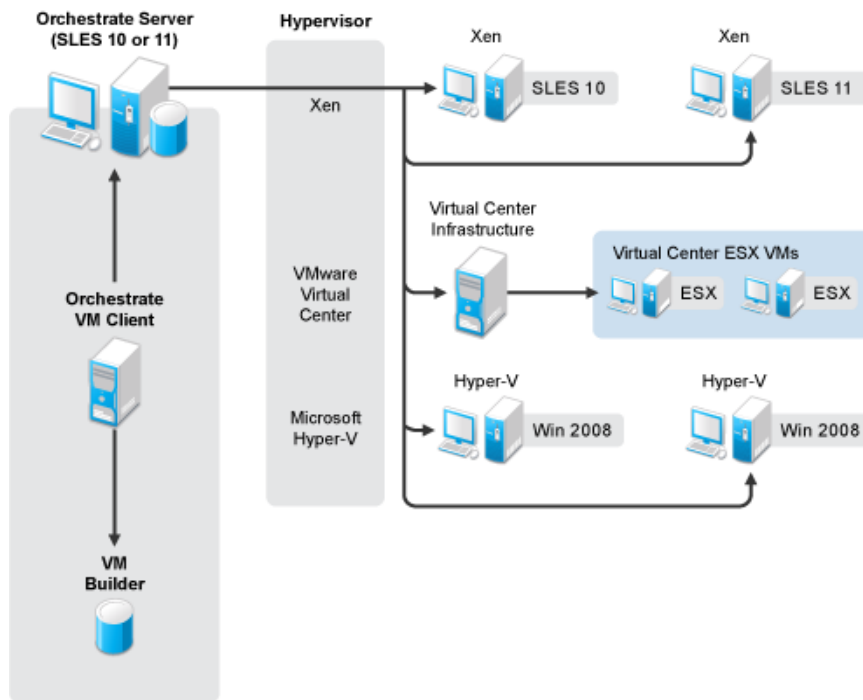
There are two types of VMs that can be provisioned:

- ♦ **VM Instance:** A VM instance is a VM that is “state-full.” This means there can only ever be one VM that can be provisioned, moved around the infrastructure, and then shut down, yet maintains its state.
- ♦ **VM Template:** A VM template represents an image that can be cloned. After it is finished its services, it is shut down and destroyed.

It can be thought of as a “golden master.” The number of times a golden master or template can be provisioned or cloned is controlled through constraints that you specify when you create a provisioning job.

The following graphic is a representation of the provisioning adapters and the way they function to communicate joblets to VMs:

Figure 9-2 VM Management Provisioning Communications



NOTE: The Xen VM Monitor can support more than just SUSE Linux Enterprise (SLE) 10 (which uses Xen 3.0.4) and Red Hat Enterprise Linux (RHEL) 5 (which uses Xen 3.0.3) VMs. For a complete list of supported guest operating systems, see the [Xen Web site \(http://www.xen.org/\)](http://www.xen.org/).

NOTE: Novell considers the PlateSpin Orchestrate provisioning technology used in conjunction with VMware ESX as experimental because it has not yet been fully tested.

The following sections provide more information on provision of VMs:

- ◆ [Section 9.5.1, “Provisioning VMs Using Jobs,” on page 84](#)
- ◆ [Section 9.5.2, “VM Placement Policy,” on page 86](#)
- ◆ [Section 9.5.3, “Provisioning Example,” on page 87](#)

9.5.1 Provisioning VMs Using Jobs

The following actions can be performed by jobs:

- ◆ Provision (schedule or manually provision a set of VMs at a certain time of day).
- ◆ Move
- ◆ Clone (clone a VM, an online VM, or a template)
- ◆ Migrate
- ◆ Destroy

- ◆ Restart
- ◆ Check status
- ◆ Create a template to instance
- ◆ Create an instance to template
- ◆ Affiliate with a host
- ◆ Make it a stand-alone VM
- ◆ Create checkpoints
- ◆ Restore
- ◆ Delete
- ◆ Cancel Action.

You might want to provision a set of VMs at a certain time of day before the need arises. You also might create a job to shut down all VMs or a constrained group of VMs. You can perform these tasks programmatically (using a job), manually (through the Development Client), or automatically on demand.

When performing tasks automatically, a job might make a request for an unavailable resource, which triggers a job to look for a suitable VM image and host. If located, the image is provisioned and the instance is initially reserved for calling a job to invoke the required logic to select, place, and use the newly provisioned resource.

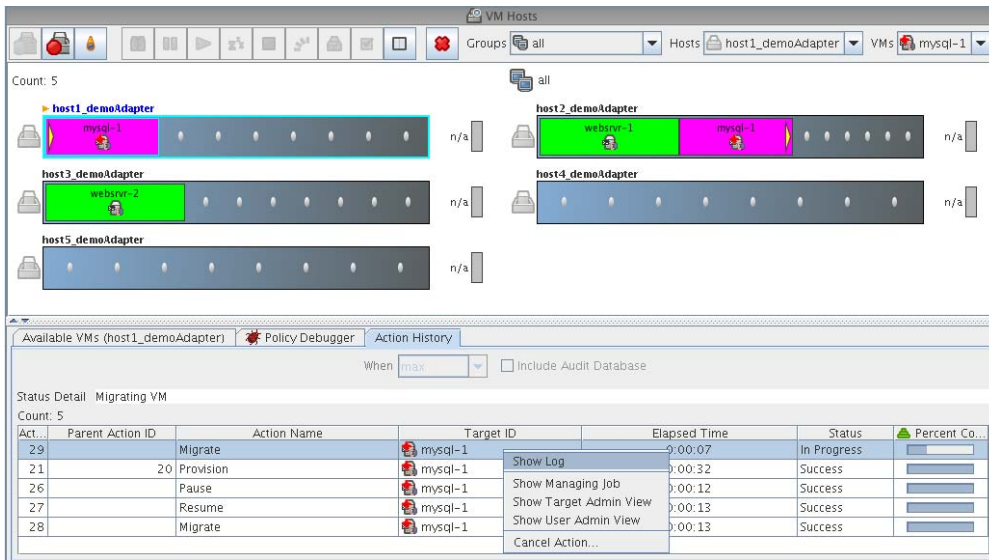
For an example of this job, see [sweeper.job \(page 155\)](#).

VM operations are available on the [ResourceInfo \(page 244\)](#) grid object, and VmHost operations are available on the [VMHostInfo \(page 250\)](#) grid object. In addition, as shown in [Section 9.5.3, “Provisioning Example,” on page 87](#), three provisioner events are fired when a provision action has completed, failed, or cancelled.

The API is equivalent to the actions available within the Development Client. The selection and placement of the VM host is governed by policies, priorities, queues, and ranking, similar to the processes used selecting resources.

Provisioning adapters on the Orchestrate Server abstract the VM. These adapters are special provisioning jobs that perform operations for each integration with different VM technologies. The following figure shows the VM host management interface that is using the Development Client.

Figure 9-3 VM Hosts Management



9.5.2 VM Placement Policy

To provision virtual machines, a suitable host must be found. The following shows an example of a VM placement policy:

```
<policy>
  <constraint type="vmhost">
    <and>
      <eq fact="vmhost.enabled" value="true"
        reason="VmHost is not enabled" />
      <eq fact="vmhost.online" value="true"
        reason="VmHost is not online" />
      <eq fact="vmhost.shuttingdown" value="false"
        reason="VmHost is shutting down" />
      <lt fact="vmhost.vm.count" factvalue="vmhost.maxvmslots"
        reason="VmHost has reached maximum vmslots" />
      <ge fact="vmhost.virtualmemory.available"
        factvalue="resource.vmimage.virtualmemory"
        reason="VmHost has insufficient virtual memory for guest VM" />
      <contains fact="vmhost.vm.availableids"
        factvalue="resource.id"
        reason="VmImage is not available on this VmHost" />
    </and>
  </constraint>
</policy>
```

9.5.3 Provisioning Example

This job example provisions a virtual machine and monitors whether provisioning completed successfully. The VM name is “webserver” and the job requires a VM to be discovered before it is run. After the provision has started, one of the three provisioner events is called.

```
1 class provision(Job):
2
3     def job_started_event(self):
4         vm = getMatrix().getGridObject(TYPE_RESOURCE, "webserver")
5         vm.provision()
6         self.setFact("job.autoterminate", False)
7
8     def provisioner_completed_event(self, params):
9         print "provision completed successfully"
10        self.setFact("job.autoterminate", True)
11
12    def provisioner_failed_event(self, params):
13        print "provision failed"
14        self.setFact("job.autoterminate", True)
15
16    def provisioner_cancelled_event(self, params):
17        print "provision cancelled"
18        self.setFact("job.autoterminate", True)
```

See additional provisioning examples in [Section 9.4, “Manual Management of a VM Lifecycle,” on page 80](#) and [Section 9.6, “Automatically Provisioning a VM,” on page 87](#).

9.6 Automatically Provisioning a VM

If you write jobs to automatically provision virtual machines, you set the following facts in the job policy:

```
resource.provision.maxcount
resource.provision.maxpending
resource.provision.hostselection
resource.provision.maxnodefailures
resource.provision.rankby
```

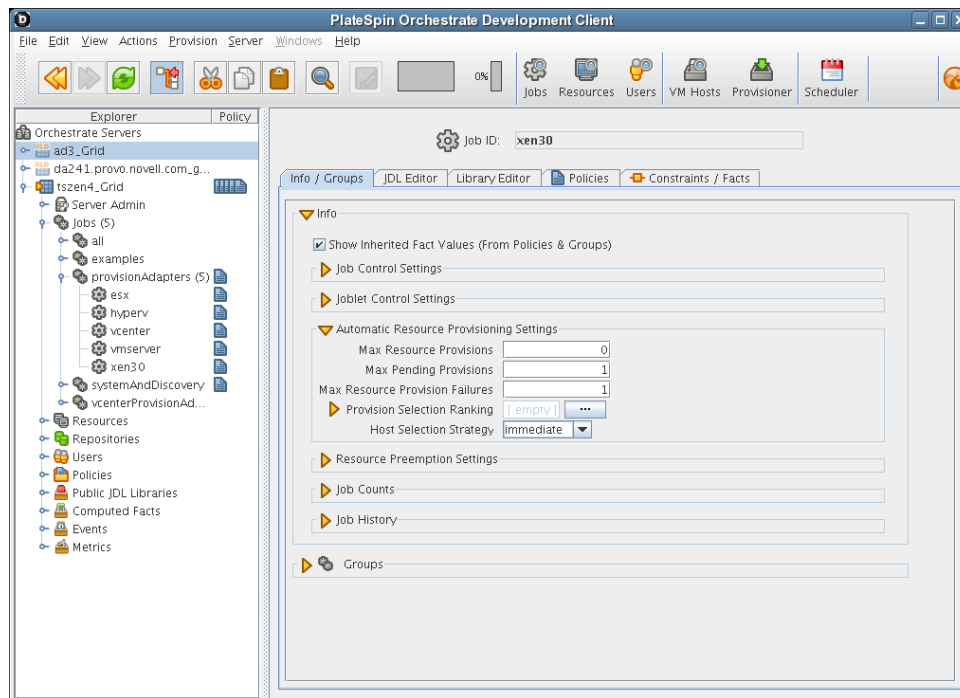
These are the job facts to enable and configure the usage of virtual machines for resource allocation. These facts can be set in a job’s policy.

For example, setting the provision.maxcount to greater than 0 allows for virtual machines to be included in resource allocation:

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

The following figure shows the job's Development Client settings to use VMs:

Figure 9-4 Job Settings for VM Provisioning



When using automatic provisioning, the provisioned resource is reserved for the job requesting the resource. This prevents another job requiring resources from obtaining the provisioned resource.

When the job that reserved the resource has finished its work (joblet has completed) on the provisioned resource, then the reservation is relaxed allowing other jobs to use the provisioned resource.

Using JDL, the reservation can be specified to reserve by JobID and also user. This is done using the [ProvisionSpec \(page 242\)](#) class.

9.7 Defining Values for Grid Objects

The following sections describe the PlateSpin Orchestrate Server grid objects and facts that are required for provisioning of PlateSpin Orchestrate resource objects. This section highlights the facts that are expected to be set from a virtual machine discovery.

- ◆ [Section 9.7.1, “PlateSpin Orchestrate Grid Objects,” on page 89](#)
- ◆ [Section 9.7.2, “Repository Objects and Facts,” on page 90](#)
- ◆ [Section 9.7.3, “VmHost Objects and Facts,” on page 96](#)
- ◆ [Section 9.7.4, “VM Resource Objects and Other Base Resource Facts,” on page 101](#)
- ◆ [Section 9.7.5, “Physical Resource Objects and Additional Facts,” on page 108](#)

9.7.1 PlateSpin Orchestrate Grid Objects

The following table explains the abbreviated codes used to define the PlateSpin Orchestrate grid objects and facts listed in the following sections:

Table 9-1 *PlateSpin Orchestrate Grid Object Definitions*

Value	Description
Automatic	The fact should be automatically set after the successful discovery of virtual resources (VmHosts and VMs).
Boolean	The fact is a Boolean value.
Default	The specified default value of the fact is set.
Dictionary	The fact is selected from a specified dictionary listing.
Dynamic	The fact is dynamically generated.
Enumerate	The fact is a specified enumerated value.
Example	When available, provides an example how a fact might be applied to an object.
Integer	The fact is an integer value.
Real	The fact is a real number.
String	The fact is a string value.
Datagrid	Facts relate to datagrid object types.
Local	Facts relate to local object types.
NAS	Facts relate to Network Attached Storage (NAS) object types.
SAN	Facts relate to Storage Area Network (SAN) object types.
Virtual	Facts relate to virtual object types.

9.7.2 Repository Objects and Facts

Facts marked with an X designate that they should be automatically set after the successful discovery of virtual resources (VmHosts and VMs). Unless marked with the ° symbol, all of the following repository objects and facts must be set for the particular provisioning adapter to function. Facts marked with °° indicate the fact is required under certain conditions.

Table 9-2 *Repository Objects and Facts*

Fact Name	Description	Fact Type	Type: X = automatically set ° = Not necessary to be set °° = Required under certain conditions
repository.capacity	The maximum amount of storage space available to virtual machines (in megabytes). The value -1 means unlimited.	Integer	<p>Local: Note: Not auto discovered, but set to a default value of -1 (unlimited size). The Administrator should alter this value.</p> <p>This fact is not currently applicable to SAN because you cannot move file-based disks into a SAN.</p> <p>SAN: Note: Not auto discovered, but set to a default value of -1 (unlimited size). The Administrator should alter this value.</p> <p>nas: Note: Not auto discovered, but set to a default value of -1 (unlimited size). The Administrator should alter this value.</p> <p>datagrid: Note: Not auto discovered, but set to a default value of -1 (unlimited size). The Administrator should alter this value.</p> <p>virtual: Note: Not auto discovered, but set to a default value of -1 (unlimited size). The Administrator should alter this value.</p>

Fact Name	Description	Fact Type	Type: X = automatically set ^o = Not necessary to be set ^{oo} = Required under certain conditions
repository.searchpath	The relative path from the location to search for VM configuration files, which implicitly includes repository.image.preferredpath.	String []	Local: X. [etc/xen/vm, myimages] <hr/> NOTE: The path is relative to repository.location or the leading '/' is ignored. <hr/> SAN: ^o . nas: X. ["my_vms", "saved_vms"] or [""] Specifies to search the whole mount. <hr/> NOTE: The path is either relative to repository.location; the leading '/' is ignored. <hr/> datagrid: N/A virtual: N/A
repository.description	The description of the repository.	String	Local: ^o Default empty. SAN: ^o . nas: ^o . datagrid: ^o Default empty. virtual: ^o Default empty.
repository.encyency	The efficiency coefficient used 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 a score. Thus, thus 0 means no cost and is very efficient).	Real	Local: Defaults to 1, which normalizes the transfer efficiency for moving VM disks. SAN: ^o Defaults to 1, which normalizes the transfer efficiency for moving VM disks. Not currently applicable because file-based disks cannot be moved into a SAN. nas: Defaults to 1, which normalizes the transfer efficiency for moving VM disks. datagrid: Defaults to 1, which normalizes the transfer efficiency for moving VM disks. virtual: Defaults to 1, which normalizes the transfer efficiency for moving VM disks.

Fact Name	Description	Fact Type	Type: X = automatically set ° = Not necessary to be set °° = Required under certain conditions
repository.enabled	True if the Repository is enabled, meaning that new VM instances can be provisioned.	Boolean	Local: Defaults to true. SAN: Defaults to true. nas: Defaults to true. datagrid: Defaults to true. virtual: Defaults to true.
repository.freespace	The amount of storage space available to new virtual machines (in megabytes). The value -1 means unlimited.	Integer	Local: Dynamic: (capacity—used space) or -1 if capacity is unlimited. SAN: Dynamic: (capacity—used space) or -1 if capacity is unlimited. nas: Dynamic: (capacity—used space) or -1 if capacity is unlimited. datagrid: Dynamic: (capacity—used space) or -1 if capacity is unlimited. virtual: Dynamic: (capacity—used space) or -1 if capacity is unlimited.
repository.groups	The groups this Repository is a member of.	String[]	Local: X SAN: X nas: X virtual: X
repository.id	The repository's unique name.	String	Local: X SAN: X nas: X datagrid: X. Currently one datagrid repository is supported. virtual: X

Fact Name	Description	Fact Type	Type: X = automatically set ° = Not necessary to be set °° = Required under certain conditions
repository.preferredpath	The relative path from the location to search and place VM files for movement and cloning.	String	Local: X. "var/lib/xen/images" <hr/> NOTE: The path is relative to repository.location; the leading '/' is ignored. <hr/> SAN: nas: X. "my_vms" <hr/> NOTE: The path is relative to repository.location; the leading '/' is ignored. <hr/> datagrid: N/A virtual: N/A
repository.location	The Repository's physical location.	String	Local: X. "/" or /var/xen/images. SAN: °. nas: X. /u or /mnt/myshareddisk. <hr/> NOTE: This is the "mount point," which is assumed to be the same mount point on every host that has a connection to this NAS. <hr/> datagrid: °X. grid:///vms virtual: N/A
repository.provisioner.jobs	The names of the provisioning adapter jobs that can manage VMs on this repository.	String []	Local: X. ["xen30"] SAN: nas: X. ["xen30"] datagrid: X. ["xen30"] virtual: X. ["vcenter"]
repository.san.type	The type of SAN (Adapter specific, "iscsi", or "fibrechannel" or " if not applicable.	String (enum)	Local: N/A, empty. SAN: Administrator must set to "iqn", "npiv", or "emc." nas: N/A, empty. datagrid: N/A, empty. virtual: N/A, empty.

Fact Name	Description	Fact Type	Type: X = automatically set ° = Not necessary to be set °° = Required under certain conditions
repository.san.vendor	The vendor of the SAN. Controls which storage bind logic to run (e.g. LUN masking, etc.).	String	Local: N/A, empty. SAN: Administrator must set to "iscsi" or "fibrechannel." nas: N/A, empty. datagrid: N/A, empty. virtual: N/A, empty.
repository.type	The type of repository: <ul style="list-style-type: none"> ◆ Local; e.g. a local disk. ◆ nas; e.g. a NFS mount. ◆ san , datagrid: A PlateSpin Orchestrate built in datagrid backed store. ◆ virtual: An externally managed VM; e.g. VMWare Virtual Center. 	String (enum)	Local: X. Local SAN: nas: datagrid: X. Datagrid virtual: X. Virtual

Fact Name	Description	Fact Type	Type: X = automatically set ° = Not necessary to be set °° = Required under certain conditions
repository.usedspace	The amount of storage space used for virtual machines.	Integer	<p>Local: Dynamic: Sum of disk space used by contained VMs. Only includes disks that are stored as local files (not partitions).</p> <p>SAN: Dynamic: Sum of disk space used by contained VMs. Only includes disks that are stored as local files (not partitions).</p> <p>Not currently applicable to SAN because you cannot move file-based disks into SAN.</p> <p>nas: Dynamic: Sum of disk space used by contained VMs. Only includes disks that are stored as local files (not partitions).</p> <p>datagrid: Dynamic: Sum of disk space used by contained VMs. Only includes disks that are stored as local files (not partitions).</p> <p>virtual: Dynamic: Sum of disk space used by contained VMs. Only includes disks that are stored as local files (not partitions).</p>
repository.vmhosts	The list of VM hosts capable of using this repository (aggregated from the individual VM host fact).	String []	<p>Local: X</p> <p>SAN:</p> <p>nas: X</p> <p>datagrid: X</p> <p>virtual: X</p>
repository.vmimages	The list of VM images stored in this repository (aggregated from individual VM fact).	String []	<p>Local: X</p> <p>SAN:</p> <p>nas: X</p> <p>datagrid: X</p> <p>virtual: X</p>

9.7.3 VmHost Objects and Facts

Unless marked with a “°” symbol, all of the following VmHost objects and facts must be set for the particular provisioning adapter to function. The “X” mark designates that the fact should be automatically set after the successful discovery of virtual resources (VmHosts and VMs).

Table 9-3 *VmHost Objects and Facts*

Fact Name	Description	Fact Type	Provision Adapter
			X = automatically set ° = Not necessary to be set °° = Required under certain conditions
vmhost.accountinggroup	The default vmhost (resource) group which is adjusted for VM statistics.	String	xen30: X. All. vmserver: X. All. vcenter: X. All.
vmhost.enabled	True if the VM host is enabled, which enables new VM instances to be provisioned.	Boolean	xen30: X. True. vmserver: X. True. vcenter: X. True.
vmhost.groups	The groups this VM host is a member of. Alias for 'vmhost.resource.group'.	String []	xen30: X. vmserver: X. vcenter: X.
vmhost.id	The VM host's unique name.	String	xen30: X. <physical host id>_xen30 vmserver: X. <physical host id>_vmserver vcenter: X. <physical host id>_vcenter
vmhost.loadindex.slots	The loading index; the ratio of active hosted VMs to the specified maximum.	Dynamic Real	xen30: X. vmserver: X. vcenter: X.
vmhost.loadindex.virtualmemory	The loading index; the ratio of consumed memory to the specified maximum.	Dynamic Real	xen30: X. vmserver: X. vcenter: X.
vmhost.location	The VM host's physical location.	String	xen30: °Defaults to empty string. vmserver: °Defaults to empty string. vcenter: X. Virtual center's 'locator' to the Vmhost; e.g., "/vcenter/eng/esx1".

Fact Name	Description	Fact Type	Provision Adapter
			X = automatically set ° = Not necessary to be set ∞ = Required under certain conditions
vmhost.maxvmslots	The maximum number of hosted VM instances.	Integer	xen30: Defaults to 3. Should be reset by Administrator. vmserver: Defaults to 3. Should be reset by Administrator. vcenter: Defaults to 3. Should be reset by Administrator.
vmhost.memory.available	The amount of memory available to new virtual machines.	Dynamic Integer	xen30: X. Calculated to be 'vmhost.memory.max; the memory consumed by running VMs. vmserver: X. Calculated to be 'vmhost.memory.max; the memory consumed by running VMs. vcenter: X. Calculated to be 'vmhost.memory.max; the memory consumed by running VMs.
vmhost.memory.max	The maximum amount of memory available to virtual machines (in megabytes).	Integer	xen30: X. Discovered. vmserver: X. Discovered. vcenter: X. Discovered.
vmhost.migration	True if the VM host can support VM migration; also subject to provision adapter capabilities.	Boolean	xen30: X. Defaults to false. Not discovered. Administrator should enable as appropriate to indicate that the VmHost supports migration. vmserver: X. Defaults to false. Not discovered. Should not be set to true since vmserver/gsx does not support migration. vcenter: X. Discovered.
vmhost.provisioner.job	The name of the provisioning adapter job that manages VM discovery on this host.	String	xen30: X. xen30. vmserver: X. vmserver. vcenter: X. vcenter.
vmhost.provisioner.password	The password required for provisioning on the VM host. This fact is used by the provisioning adapter.	String	xen30: °. vmserver: °. If set, this fact is passed to the vmserver CLI tools to authenticate. Not necessary if Orchestrate Agent is run as root. vcenter: °.

			Provision Adapter
Fact Name	Description	Fact Type	X = automatically set ^o = Not necessary to be set ^{oo} = Required under certain conditions
vmhost.provisioner.username	The username required for provisioning on the VM host. This fact is used by the provisioning adapter.	String	xen30: ^o . vmserver: If set, is passed to vmserver CLI tools to authenticate. Not necessary if Orchestrate Agent is run as root. vcenter: ^o .
vmhost.repositories	This list of repositories (VM disk stores) is visible to this VM host.	String []	xen30: X. Discovery only adds the local repository and the datagrid on the first creation of the vmhost. Administrator is required to add SAN/NAS repositories or remove local if desired. vmserver: X. Discovery only adds the local repository and the datagrid on the first creation of the vmhost. Administrator is required to add SAN/NAS repositories or remove local if desired. vcenter: X. Automatically set to VirtualCenter. <hr/> NOTE: This is the only sensible setting.
vmhost.resource	The name of the resource that houses this VM host container.		xen30: ^o . vmserver: X. vcenter: X.
vmhost.shuttingdown	True if the VM host is attempting to shut down and does not need to be provisioned.	Dynamic Boolean	xen30: Initially False, then set to True when the administrator specifies to shut down a host. vmserver: Initially False, then set to True when the administrator specifies to shut down a host. vcenter: Initially False, then set to True when the administrator specifies to shut down a host.

Fact Name	Description	Fact Type	Provision Adapter
			X = automatically set ° = Not necessary to be set °° = Required under certain conditions
vmhost.vm.available.groups	The list of resource groups containing VMs that are allowed to run on this host.	String []	<p>xen30: X. Automatically set to VMs_<provisioning_adapter>; e.g., any VM of compatible type can be provisioned. The VMs_<provisioning_adapter> group is automatically created by discovery. The administrator can refine this by creating new groups and editing if further restrictions are required.</p> <p>vmserver: X. Automatically set to VMs_<provisioning_adapter>; e.g., any VM of compatible type can be provisioned. The VMs_<provisioning_adapter> group is automatically created by discovery. The administrator can refine this by creating new groups and editing if further restrictions are required.</p> <p>vcenter: X. Discovery attempts to map Virtual Center grouping to PlateSpin Orchestrate resources groups and sets this fact accordingly. This also includes a special "template_vcenter" group to map to Virtual Center 1.3.x "templates".</p>
vmhost.vm.count	The current number of active VM instances.	Dynamic Integer	<p>xen30: X.</p> <p>vmserver: X.</p> <p>vcenter: X.</p>
vmhost.vm.instanceids	The list of active VM instances.	Dynamic String[]	<p>xen30: X.</p> <p>vmserver: X.</p> <p>vcenter: X.</p>
vmhost.vm.templatecounts	A dictionary of running instance counts for each running VM template.	Dynamic Dictionary	<p>xen30: X.</p> <p>vmserver: X.</p> <p>vcenter: X.</p>
vmhost.xen.bits	xen30 only. Legal values are 32 and 64.	Integer	<p>xen30: X. 64.</p> <p>vmserver: °Not defined.</p> <p>vcenter: °Not defined.</p>

Fact Name	Description	Fact Type	Provision Adapter
			X = automatically set ° = Not necessary to be set °° = Required under certain conditions
vmhost.xen.hvm	xen30 only.	Boolean	xen30: X. True. vmserver: °Not defined. vcenter: °Not defined.
vmhost.xen.version	xen30 only: Major.Minor version of the Xen hypervisor.	Real	xen30: X. 3.00 vmserver: °Not defined. vcenter: °Not defined.
vmhost.vcenter.hostname	vcenter only. The hostname of the resource containing this VM container. NOTE: Deprecated. Use 'vmhost.resource.host name instead.	String	xen30: °Not defined. vmserver: °Not defined. vcenter: X. esx1.
vmhost.vcenter.networks	vcenter only. List of network interfaces on the physical host.	List	xen30: °Not defined. vmserver: °Not defined. vcenter: VM network.
vmhost.vcenter.grouppath	vcenter only: Part of the Virtual Center "locator" URL.	List	xen30: °Not defined. vmserver: °Not defined. vcenter: X. /vcenter/eng1.

9.7.4 VM Resource Objects and Other Base Resource Facts

The following virtual machine resource objects and additional base resource facts marked with the “•” symbol must be set for the particular provisioning adapter to function. Facts marked with “••” indicate the fact is required under certain conditions. The “X” character designates that the fact should be automatically set after the successful discovery of virtual resources (VmHosts and VMs).

Table 9-4 Resource Objects (VM only) and Additional Facts to Base Resource Facts

Fact Name	Description	Type	Provision Adapter
			X = automatically set ° = Not necessary to be set °° = Required under certain conditions
resource.provision.automatic	Signifies that this resource was cloned/ provisioned automatically and thus is shut down/ destroyed automatically as well.	Dynamic Boolean	xen30: ° . vmserver: ° . vcenter: ° .
resource.provision.autoprep.*	Fact namespace used to convey configuration information actually used to "personalize" this VM instance.	<various>	xen30: ° X. Can be set when rediscovering the state or as a result of a migration or provision action. vmserver: ° X. Can be set when rediscovering the state or as a result of a migration or provision action. vcenter: ° X. Can be set when rediscovering the state or as a result of a migration or provision action.
resource.provision.currentaction	The current management action in progress on this provisionable resource.c.	Dynamic String	xen30: ° . vmserver: ° . vcenter: ° .
resource.provision.hostwait	The time (seconds) this resource has been waiting / waited for a suitable host.	Dynamic Integer	xen30: ° . vmserver: ° . vcenter: ° .

			Provision Adapter
			X = automatically set
Fact Name	Description	Type	° = Not necessary to be set
			°° = Required under certain conditions
resource.provision.jobid	The current or last job ID that performed a provisioning action on this resource. Useful for viewing the job log.	Dynamic String	xen30: ° . vmserver: ° . vcenter: ° .
resource.provision.resync	Specifies that the provisioned resource's state needs to be re-synced with the underlying provisioning technology at the next opportunity.	Dynamic Boolean	xen30: °X. Can be set on discovery when the Orchestrate state machine mismatches the VM state. This initiates a future VM state recovery action ("Check Status"). May be set for delayed re-discovery by administrator or JDL logic. vmserver: °X. Can be set on discovery when the Orchestrate state machine mismatches the VM state. This initiates a future VM state recovery action ("Check Status"). May be set for delayed re-discovery by administrator or JDL logic. vcenter: °X. Can be set on discovery when the Orchestrate state machine mismatches the VM state. This initiates a future VM state recovery action ("Check Status"). May be set for delayed re-discovery by administrator or JDL logic.
resource.provision.state	The current state of this provisioned instance (down, suspended, up, paused) or unknown if an admin action is currently being performed.	Dynamic String (enum)	xen30: ° . vmserver: ° . vcenter: ° .
resource.provision.status	The current descriptive status of the provisioned resource.	Dynamic String	xen30: ° . vmserver: ° . vcenter: ° .

Fact Name	Description	Type	Provision Adapter
			X = automatically set ^o = Not necessary to be set ^{oo} = Required under certain conditions
resource.provision.template	The ID of the template resource that this instance was created from (if applicable).	Dynamic String	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provision.time.request	The time when the last provision (or other administrative action) request was made.	Dynamic Date	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provision.time.shutdown	The time when the resource was last shut down.	Dynamic Date	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provision.time.start	The time when the resource was last successfully provisioned.	Dynamic Date	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provision.vmhost	The ID of the host currently housing this provisioned resource.	Dynamic String	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provisionable	True if the resources is a provisionable type.	Dynamic Boolean	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provisioner.autoprep.*	Fact namespace used to convey configuration information actually used to "personalize" this VM instance.	<various>	xen30: ^o X. vmserver: ^o . vcenter: ^o .
resource.provisioner.count	The total count of operational instances and provisions in progress"	Dynamic Integer	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provisioner.debug	Controls the debug log level in the provisioner.	Boolean	xen30: ^o . vmserver: ^o . vcenter: ^o .

Fact Name	Description	Type	Provision Adapter
			X = automatically set ^o = Not necessary to be set ^{oo} = Required under certain conditions
resource.provisioner.host.maxwait	The maximum time to wait for a suitable host before timing out (in seconds, '<0' to wait indefinitely).	Integer	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provisioner.host.preferred wait	The time after which some VMhost constraints is lifted to increase the available pool by, for example, considering moving the disk image (in seconds, <0 to wait indefinitely).	Integer	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provisioner.instances	The list of id's of the instances of this template resource (if applicable).	String[]	xen30: ^o . vmserver: ^o . vcenter: ^o .
resource.provisioner.job	The name of the provisioning job that manages the life cycle of this resource.	String	xen30: X. xen30 vmserver: X. vmserver vcenter: X. vcenter
resource.provisioner.maxinstances	The maximum allowed number of instances of this provisionable resource (applicable only to templates).	Integer	xen30: X. Defaults to 1. Administrator should reset for VM templates to allow multiple clones. vmserver: X. Defaults to 1. Administrator should reset for VM templates to allow multiple clones. vcenter: X. Defaults to 1. Administrator should reset for VM templates to allow multiple clones.
resource.provisioner.recommended host	The host on which the image for this resource is associated; e.g., was suspended or is the preferred host for quick startup.	String	xen30: ^o X. vmserver: ^o X. vcenter: ^o X.

			Provision Adapter
			X = automatically set
Fact Name	Description	Type	° = Not necessary to be set
			°° = Required under certain conditions
resource.vcenter.grouppath	Locator for the Virtual Center group that the VM resides in.	String	xen30: ° Not defined. vmserver: ° Not defined. vcenter: X. /vcenter/eng
resource.vcenter.guestOS	VMWare's name for the guest OS.	String	xen30: ° Not defined. vmserver: ° Not defined. vcenter: X. winNetEnterprise.
resource.vcenter.imagepath	Locator for the VM in Virtual Center.	String	xen30: ° Not defined. vmserver: ° Not defined. vcenter: X. /vcenter/eng/windows2003ent.
resource.vm.basepath	The file system location of the VM files either absolute or relative to the 'repository.location' fact.	String	xen30: X. Example: "var/lib/xen/images/sles10". vmserver: X. For example, "/var/lib/vmware/Virtual-Machines/sles9". Location in the repository of the directory containing VM disks, configuration file and other related files. vcenter: ° N/A.
resource.vm.configfile	The location of the VM's configuration file inside of the default repository (resource.provisioner.repository).	String	xen30: X. /etc/xen/vm/sles10. vmserver: ° Not currently used. vcenter: ° N/A.
resource.vm.cpu.architecture	The required cpu architecture e.g. x86, x86_64, sparc.	String	xen30: ° X. vmserver: ° . vcenter: ° .

Fact Name	Description	Type	Provision Adapter X = automatically set ° = Not necessary to be set °° = Required under certain conditions
resource.vm.cpu.weight	The CPU weight for this VM. A value of '1.0' represents normal weighting; setting another VM to a weight of '2.0' would mean it would get twice as much cpu as this VM.	Real	xen30: ° . vmserver: ° . vcenter: ° .
resource.vm.files	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'.	Dictionary	xen30: X. { "mof": /var/lib/xen/images/sles10/mof" , "suspendcheckpoint": "/var/lib/xend/domain/checkpoint", "config": "/var/lib/xen/images/sles10/config.xen" } . vmserver: ° X { "config": "/var/lib/vmware/Virtual Machines/sles10/sles10.vmx" } . vcenter: ° N/A.
resource.vm.maxinstancespervmhost	The maximum allowed number of instances of this VM image per vmhost.	Integer	xen30: Defaults to 1. Administrator should increase if more than one instance of the same VM template is allowed to be run on one host.
resource.vm.memory	The configured virtual memory requirement of this VM image (megabytes).	Integer	xen30: X. vmserver: X. vcenter: X.
resource.vm.preventmove	Set by the administrator to prevent relocation of a VM (disk moves) even if possible.	Boolean Default: False	
resource.vm.type	The required system type of a virtual machine ('full' or 'para').	String	xen30: X. vmserver: X. vcenter: X.

Fact Name	Description	Type	Provision Adapter X = automatically set ° = Not necessary to be set °° = Required under certain conditions
resource.vm.uuid	The UUID of a virtual machine (vendor/adaptor specific).	String	xen30: X. {vm uuid} vmserver: Not currently used. vcenter: Not currently used.
resource.vm.vcpu.number	The number of virtual CPUs for this VM.	Integer	xen30: ° X. vmserver: ° . vcenter: ° .
resource.vm.vdisks	The specification of virtual disks that make up this VM. The dictionary keys are name (String), repository (String), location (String), size (Integer), fixed (Boolean).	List of Dictionaries	xen30: X. [{ "location": "/var/lib/xen/images/sles10/disk1", "moveable": True, "repository": "vmhost1" ... }]. vmserver: ° Not currently used. vcenter: ° N/A.
resource.vm.vdisksize	The total size of all the moveable virtual desks for this VM image (megabytes).	Integer	xen30: X. vmserver: X. vcenter: X.
resource.vm.vendor	The vendor of a virtual machine.	String	xen30: ° X. vmserver: ° X. vcenter: ° X.
resource.vm.version	The version number for this VM.	Integer	xen30: X. vmserver: ° N/A. vcenter: ° N/A.
resource.vm.vmhost.rankby	The ranking specification used to select suitable vm hosts. Element syntax is <fact>/<order> where order is either a (ascending) or d (descending).	String[]	xen30: Defaults to vmhost.vm.placement.score/a, vmhost.loadindex.slots/a. vmserver: Defaults to vmhost.vm.placement.score/a, vmhost.loadindex.slots/a. vcenter: Defaults to vmhost.vm.placement.score/a, vmhost.loadindex.slots/a.

Fact Name	Description	Type	Provision Adapter X = automatically set ° = Not necessary to be set °° = Required under certain conditions
resource.vnc.ip	The host IP address for a VNC session running on the resource.	String	xen30: ° X. 192.168.0.4 vmserver: Not used. vcenter: ° .
	NOTE: Technically, this fact is available on all resources both VMs and physical.		
resource.vnc.port	The port number for a VNC session running on the resource.	Integer	xen30: ° X. 5900 vmserver: Not used. vcenter: ° .
	NOTE: Technically, this fact is available on all resources both VMs and physical.		

9.7.5 Physical Resource Objects and Additional Facts

The following physical resource objects and additional base resource facts marked with the “•” symbol must be set for the particular provisioning adapter to function. The physical resources have the potential of creating VmHost containers.

Facts marked with “••” indicate the fact is required under certain conditions. The “X” character designates that the fact should be automatically set after the successful discovery of virtual resources (VmHosts and VMs).

Table 9-5 Resource Object (Physical that have the potential for VmHost containers) / Additional Facts (additional to base resource set)

Fact Name	Description	Type	Provision Adapter
			X = automatically set ° = Not necessary to be set °° = Required under certain conditions
resource.vcenter.client	vcenter only: Marks resources and Virtual Center web services client capable.	Boolean	xen30: ° Not defined. vmserver: ° Not defined. vcenter: Administrator must set through association of 'vcenter_client.policy' with appropriate resources.
resource.vmserver.cmdpath	vmserver only: Path to VMWare CLI tools.	String	xen30: ° Not defined. vmserver: X. For example, "/usr/bin/vmware-cmd" vcenter: ° Not defined.
resource.vmserver.localrepositories	vmserver only: Paths to VM storage directories.	List	xen30: ° Not defined. vmserver: For example, "/var/lib/vmware/virtual machines" vcenter: ° Not defined.
resource.vmserver.vmruntimepath	vmserver only: Full path to vmrun CLI tool.	String	xen30: ° Not defined. vmserver: X. For example, "/usr/bin/vmrun". vcenter: ° Not defined.
resource.xen	xen30 only: Xen enabled.	Boolean	xen30: X. True. vmserver: ° Not defined. vcenter: ° Not defined.
resource.xen.bits	xen30 only: (legal values are 32 and 64)	String	xen30: X. 64 bit. vmserver: ° Not defined. vcenter: ° Not defined.
resource.xen.hvm	xen30 only:	Boolean	xen30: X. True. vmserver: ° Not defined. vcenter: ° Not defined.

Fact Name	Description	Type	Provision Adapter X = automatically set ° = Not necessary to be set °° = Required under certain conditions
resource.xen.version	xen30 only: Major.Minor version of the Xen hypervisor.	Real	xen30: X. 3.00. vmserver: ° Not defined. vcenter: ° Not defined.

Complete Job Examples

10

This section describes specific Job examples that can be deployed using the PlateSpin® Orchestrate Server from Novell®. The following sections demonstrate some practical ways to use PlateSpin Orchestrate and should help you better understand how to write your own jobs:

- ♦ [Section 10.1, “Accessing Job Examples,” on page 111](#)
- ♦ [Section 10.2, “Installation and Getting Started,” on page 111](#)
- ♦ [Section 10.3, “PlateSpin Orchestrate Sample Job Summary,” on page 112](#)
- ♦ [Section 10.4, “Parallel Computing Examples,” on page 113](#)
- ♦ [Section 10.5, “General Purpose Jobs,” on page 125](#)
- ♦ [Section 10.6, “Miscellaneous Code-Only Jobs,” on page 167](#)

10.1 Accessing Job Examples

The basic examples delivered with PlateSpin Orchestrate are located in either of two possible installation directories depending on the type of installation. For server installations, look here:

```
/opt/novell/zenworks/zos/server/examples/
```

For client installation, look here:

```
/opt/novell/zenworks/zos/client/examples/
```

When you unjar or unzip examples from the from the `<path>/examples/<example>.job` file or view jobs using the details panel and the JDL and Policy tabs in PlateSpin Orchestrate Development Client, you should see the `.jdl` and `.policy` files.

Policy files specify how the job arguments and static attributes are defined. Or, you can use the `zos jobinfo` command to simply display job arguments and their default values.

All of the examples can be opened and modified using a standard code editor, then redeployed and examined using the procedure explained in “[Walkthrough: Deploying a Sample Job](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

10.2 Installation and Getting Started

To run the PlateSpin Orchestrate described in this section, use the following guidelines:

- ♦ Install and configure PlateSpin Orchestrate properly (see “[Installing and Configuring All PlateSpin Orchestrate Components Together](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*).
- ♦ Unless otherwise indicated, install at least one agent on a managed resource and have it running (see “[Installing the PlateSpin Orchestrate Agent on Other Supported Operating Systems](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*).
- ♦ Before running `zosadmin` or `zos` commands, you must log into the Orchestrate Server.

The `zosadmin` command is required for administrating jobs. This includes deploying and undeploying a job to the server. The `zos` command is for job control, including starting a job and viewing a job's log. As you learn about the PlateSpin Orchestrate job samples, you will use the `zosadmin` command for deploying a sample job and the `zos` command for running the sample.

- ◆ For an explanation of the `zosadmin` commands, see “[The zosadmin Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*.

```
> zosadmin login --user zosadmin
Login to server: skate
Please enter current password for 'zosadmin':
Logged into grid on server 'skate'
```

- ◆ For an explanation of `zos` commands, see “[The zos Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*.

```
> zos login --user zenuser
Please enter current password for 'zenuser':
Logged into grid as zenuser
```

You could create a user (see “[Walkthrough: Creating a User Account](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*) but `zos login --user=vmmanger` works with the account created by default during installation. The password for this user was also created during the installation; that is, when either `config` or `guiconfig` was run. If you are not the person who created this password during configuration, you will need to contact the individual who did so to get the password.

10.3 PlateSpin Orchestrate Sample Job Summary

The following table provides a high-level explanation of the PlateSpin Orchestrate job examples that are delivered with PlateSpin Orchestrate and the job developer concepts you might want to understand:

Table 10-1 *PlateSpin Orchestrate Job Development Examples*

Example Name	Job Function Capabilities
demolterator.job (page 114)	<ul style="list-style-type: none"> ◆ Using policy constraints and job arguments to restrict joblet execution to specific resources. ◆ Scheduling joblets using a <code>ParameterSpace</code>. ◆ Provides an example of executing a command on a resource.
dgtest.job (page 126)	<ul style="list-style-type: none"> ◆ Downloading files stored on grid management servers to networked nodes.
factJunction.job (page 168)	<ul style="list-style-type: none"> ◆ Retrieving information about objects in the grid relative to another object.
failover.job (page 136)	<ul style="list-style-type: none"> ◆ Managing how joblets failover to enhance the robustness of your jobs.
instclients.job (page 143)	<ul style="list-style-type: none"> ◆ Installing an PlateSpin Orchestrate client on multiple machines. ◆ Provides an example of executing a command on a resource.

Example Name	Job Function Capabilities
jobargs.job (page 177)	<ul style="list-style-type: none"> Understanding the various argument types that jobs can accept (integer, real, Boolean, string, time, date, list, dictionary, and array, which can contain the types integer, real, Boolean, time, date, and String).
notepad.job (page 150)	<ul style="list-style-type: none"> Understanding how to launch specific applications on specified resources.
quickie.job (page 121)	<ul style="list-style-type: none"> Understanding how jobs can start multiple instances of a joblet on one or more resources.
sweeper.job (page 155)	<ul style="list-style-type: none"> Understanding how poll all resources on the grid.an ordered serialized scheduling of the joblets
whoami.job (page 162)	<ul style="list-style-type: none"> Sending a command to the operating system's default command interpreter. On Microsoft Windows, this is <code>cmd.exe</code>. On POSIX systems, this is <code>/bin/sh</code>.

10.4 Parallel Computing Examples

The following examples demonstrate high performance or parallel computing concepts:

- [“demoIterator.job” on page 114](#)
- [“quickie.job” on page 121](#)

demolterator.job

Reference implementation for a simple test iterator. Several concepts are demonstrated: 1) Using policy constraints and job arguments to restrict joblet execution to a specific resource, 2) Scheduling joblets using a ParameterSpace, and 3) An example of executing a command on a resource.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
Logged into grid as zenuser

> zos jobinfo --detail demoIterator
Jobname/Parameters  Attributes
-----
demoIterator        Desc: This example job is a reference for a simple test
                    iterator. It is useful for demonstrating how policies
                    and job args can be used to target the job to a
                    particular resource.

                    numJoblets      Desc: joblets to run
                    Type: Integer
                    Default: 100

                    cmd              Desc: Simple command to execute
                    Type: String          Default:

                    os              Desc: Regular expression match for Operating System
                    Type: String
                    Default: .*

                    cpu             Desc: Regular expression match for CPU architecture
                    Type: String
                    Default: .*
```

Description

The files that make up the DemoIterator job include:

```
demoIterator          # Total: 156 lines
|-- demoIterator.jdl  #   79 lines
`-- demoIterator.policy #   77 lines
```

demolterator.jdl

```
1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
```

```

CONTRACT,
 9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: demoIterator.jdl,v 1.4 2008/03/05 20:05:48 ray Exp $
13 # -----
-----
14
15 import time, random
16
17 #
18 # Add to the 'examples' group on deployment
19 #
20 if __mode__ == "deploy":
21     try:
22         jobgroupname = "examples"
23         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
24         if jobgroup == None:
25             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
26         jobgroup.addMember(__jobname__)
27     except:
28         exc_type, exc_value, exc_traceback = sys.exc_info()
29         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
30
31
32 class demoIteratorJob(Job):
33
34     def job_started_event(self):
35         print 'job_started_event'
36         self.completed = 0
37
38         # Launch the joblets
39         numJoblets = self.getFact("jobargs.numJoblets")
40         print 'Launching ', numJoblets, ' joblets'
41
42         pspace = ParameterSpace()
43         i = 1
44         while i <= numJoblets:
45             pspace.appendRow({'name':'joblet'+str(i)})
46             i += 1
47         pspace.maxJobletSize = 1
48         self.schedule(demoIteratorJoblet,pspace,{})
49
50     def joblet_completed_event(self, jobletnumber, node):
51         self.completed += 1
52         self.setFact("jobinstance.memo", "Tests run: %s" %
(self.completed))
53
54
55 class demoIteratorJoblet(Joblet):
56
57     def joblet_started_event(self):
58         print "Hi from joblet ", self.getFact("joblet.number")
59         time.sleep(random.random() * 15)
60
61         cmd = self.getFact("jobargs.cmd")

```

```

62         if len(cmd) > 0:
63             system(cmd)
64
65
66
67             # Example of more sophisticated exec
68             # e.g. e.signal("SIGUSR1")
69             """
70             e = Exec()
71             e.setCommand(cmd)
72             #e.setStdoutFile("cmd.out")
73             e.writeStdoutToLog()
74             e.writeStderrToLog()
75             #try:
76             e.execute()
77             #except:
78                 #self.retry("retryable example error")
79             """

```

demolterator.policy

```

1  <!--
2
*=====
=
3  * Copyright (c) 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT. NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
*=====
=
13 * $Id: demoIterator.policy,v 1.2 2008/02/27 20:49:34 john Exp $
14
*=====
=
15 -->
16
17 <policy>
18   <constraint type="accept" reason="Too busy for more work. Try again
later!">
19     <or>
20       <lt fact="job.instances.queued" value="4" />
21       <contains fact="user.groups" value="superuser" />
22     </or>
23   </constraint>
24
25   <constraint type="start" reason="Waiting on queue">
26     <or>

```

```

27     <lt fact="job.instances.active" value="2" />
28     <contains fact="user.groups" value="superuser" />
29   </or>
30 </constraint>
31
32 <jobargs>
33   <fact name="numJoblets"
34     type="Integer"
35     description="joblets to run"
36     value="100"
37     visible="true" />
38
39   <fact name="cmd"
40     type="String"
41     description="Simple command to execute"
42     value="" />
43
44   <fact name="os"
45     type="String"
46     description="Regular expression match for Operating System"
47     value=".*" />
48
49   <fact name="cpu"
50     type="String"
51     description="Regular expression match for CPU architecture"
52     value=".*" />
53 </jobargs>
54
55 <constraint type="resource" reason="Does not match">
56   <and>
57     <eq fact="resource.os.family" factvalue="jobargs.os" match="regexp"
58     />
59     <eq fact="resource.cpu.architecture" factvalue="jobargs.cpu"
60     match="regexp"/>
61   </or>
62   <and>
63     <defined fact="env.VENDOR" />
64     <eq fact="resource.os.vendor" factvalue="env.VENDOR"
65     match="regexp" />
66   </and>
67   <undefined fact="env.VENDOR" />
68 </or>
69 </constraint>
70
71 <job>
72   <fact name="description"
73     type="String"
74     value="This example job is a reference for a simple test
75     iterator. It is useful for demonstrating how policies and job args can be used
76     to target the job to a particular resource." />
77 </job>

```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Exec

Used to manage command line execution on resources.

ParameterSpace

Defines a parameter space to be used by the scheduler to create a Joblet set. A parameter space might consist of rows of columns or a list of columns that is expanded and can be turned into a cross product.

Job Details

The following sections describe the DemoIterator job:

- ♦ [“zosadmin deploy” on page 118](#)
- ♦ [“job_started_event” on page 119](#)
- ♦ [“joblet_started_event” on page 119](#)

zosadmin deploy

The deployment for the DemoIterator job is performed by lines 20-29 of [demoIterator.jdl \(page 114\)](#). When jobs are deployed into the grid, they can optionally be organized for grouping. In this case, the demoIterator job is added to the group named examples, and can be displayed in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples
```

For a general overview of how jobs are added to groups during deployment, see [“Walkthrough: Deploying a Sample Job”](#) in the *“PlateSpin Orchestrate 2.0 Installation and Configuration Guide.”*

job_started_event

When the DemoIterator job receives a `job_started_event`, it creates a `ParameterSpace` JDL class and adds the number of rows as indicated by the value of the argument `numJoblets` (see lines 42-46 in [demoIterator.jdl \(page 114\)](#)). A `ParameterSpace` object is like a spreadsheet, containing rows and columns of information that might all be given to one joblet or sliced up across many joblets at schedule time. In this case, the `ParameterSpace` is told that `maxJobletSize` is 1 (see line 47), meaning a joblet instance is created for each row in the `ParameterSpace` during job scheduling (see line 48).

Not shown in this example is the fact that a joblet can get access to this “spreadsheet” of information by calling `self.getParameterSpace()`, and calling `hasNext()` and `next()` to enumerate through each row of information. To learn more about putting information in a `ParameterSpace` object from a job and obtaining that information from the `JobletParameterSpace` object from a joblet, see [ParameterSpace \(page 240\)](#).

The resource that runs the joblet is determined from the resource constraint specified in lines 18-30 and 55-68 of [demoIterator.policy \(page 116\)](#), and from the values specified for the parameters `os` and `cpu` supplied on the command line. If these parameters are not specified on the command line, the default value for both is the regular expression `.*`, which means to include everything.

The constraints at lines 18-30 in [demoIterator.policy \(page 116\)](#) define the work load for the resources. In this case, resources do not accept jobs if there are already four jobs queued up, and are not to run jobs if there are two or more jobs currently in progress.

To learn more about setting `start`, `resource`, or `accept` constraints in a policy file, see [“Defining Job Elements” on page 56](#).

joblet_started_event

As the DemoIterator joblet is executed on a particular resource, it receives a `joblet_started_event`. When this happens, the DemoIterator joblet simply sleeps for a random amount of time to stagger the execution of the joblets and then sends a command to the operating system, if one was supplied as a job argument. The command is executed on the target operating system using the built-in function `system()`, which is an alternative to using the more feature-rich class `Exec`.

For more information on sending commands to the operating system using the `Exec` class, see [Exec](#).

After the joblet is finished running, a `joblet_completed_event` is sent to `demoIteratorJob`, which increments the variable `completed`, and posts the updated value to the job fact `jobinstance.memo` (see lines 50-52 in [demoIterator.jdl \(page 114\)](#)). You can see the text for the memo displayed on the Job Log tab in the list of running jobs in the PlateSpin Orchestrate Development Client.

For more information, see [“Starting and Stopping the PlateSpin Orchestrate Development Client”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

Configure and Run

Execute the following commands to deploy and run `demoIterator.job`:

- 1 Deploy `demoIterator.job` into the grid:

```
> zosadmin deploy demoIterator.job
```
- 2 Display the list of deployed jobs:

```
> zos joblist
```

demoIterator should appear in this list.

- 3 Run the job on the first available resource without regard to OS or CPU, and use the default value for number of joblets, which is 100:

```
> zos run demoIterator
```

- 4 Run 10 joblets on Intel Windows resources, and launch the Notepad* application on each one:

```
> zos run demoIterator numJoblets=10 cmd=notepad os=Windows cpu=i386
```

NOTE: If a resource with the matching OS is not available, the job remains in the “waiting” state.

Here is an example that runs the `pwd` command on three joblets on the Linux operating system:

```
> zos run demoIterator numJoblets=3 cmd=pwd os=linux
JobID: zenuser.demoIterator.417
```

```
zos log zenuser.demoIterator.417
job_started_event
Launching 3 joblets
[freeze] Hi from joblet 1
[freeze] /var/opt/novell/zenworks/zos/agent/node.default/freeze/
zenuser.demoIterator.417.1
[skate] Hi from joblet 0
[skate] /var/opt/novell/zenworks/zos/agent/node.default/skate/
zenuser.demoIterator.417.0
[melt] Hi from joblet 2
[melt] /var/opt/novell/zenworks/zos/agent/node.default/melt/
zenuser.demoIterator.417.2
```

See Also

- ♦ Setting Constraints Using Policies (see [Section 4.4, “Policy Management,”](#) on page 34 and [Chapter 5, “Policy Elements,”](#) on page 45).
- ♦ Adding Jobs to Groups During Deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job](#) (page 168)).
- ♦ [quickie.job](#) (page 121) demonstrates how a job starts up multiple instances of a joblet on one or more resources. The `Joblet` class defines how a joblet is executed on a resource.
- ♦ Setting default parameter values using policies
- ♦ Configuring constraints in a policy file
- ♦ Naming conventions for policy facts (see [Section 3.1.1, “PlateSpin Orchestrate Datagrid Filepaths,”](#) on page 25)
- ♦ Facts provided by the PlateSpin Orchestrate system that can be referenced within a JDL file
- ♦ Using the `zos` command line tool (see [“The zos Command Line Tool”](#) in the *PlateSpin Orchestrate 2.0 Command Line Reference*).
- ♦ Running commands using the `Exec` class

quickie.job

Demonstrates a job starting up multiple instances of a joblet on one or more resources. Because this job simply launches and returns immediately, it can also be useful for testing network latency.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
Logged into grid as zenuser

> zos jobinfo --detail quickie
Jobname/Parameters  Attributes
-----
quickie              Desc: This example job does absolutely nothing. It just
                    returns immediately. For testing network latency.

                    sleeptime      Desc: time to sleep (in seconds)
                    Type: Integer
                    Default: 0

                    numJoblets     Desc: joblets to run
                    Type: Integer
                    Default: 100
```

Description

The files that make up the Quickie job include:

```
quickie              # Total: 88 lines
|-- quickie.jdl      # 48 lines
`-- quickie.policy   # 40 lines
```

quickie.jdl

```
1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: quickie.jdl,v 1.3 2008/02/27 20:51:13 john Exp $
13 # -----
-----
14
```

```

15 import time
16
17 #
18 # Add to the 'examples' group on deployment
19 #
20 if __mode__ == "deploy":
21     try:
22         jobgroupname = "examples"
23         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
24         if jobgroup == None:
25             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
26         jobgroup.addMember(__jobname__)
27     except:
28         exc_type, exc_value, exc_traceback = sys.exc_info()
29         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
30
31
32 class quickieJob(Job):
33
34     def job_started_event(self):
35
36         # Launch the joblets
37         numJoblets = self.getFact("jobargs.numJoblets")
38         print 'Launching ', numJoblets, ' joblets'
39
40         self.schedule(quickieJoblet, numJoblets)
41
42
43 class quickieJoblet(Joblet):
44
45     def joblet_started_event(self):
46         self.setFact("joblet.memo", "quickie's memo - joblet started")
47         sleeptime = self.getFact("jobargs.sleeptime")
48         time.sleep(sleeptime)

```

quickie.policy

```

1  <!--
2
*=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT.  NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
*=====
=

```

```

13 * $Id: quickie.policy,v 1.2 2008/02/27 20:51:13 john Exp $
14
15 =====
16 =
17 -->
18 <policy>
19     <jobargs>
20         <fact name="numJoblets"
21             type="Integer"
22             description="joblets to run"
23             value="100"
24             visible="true" />
25
26         <fact name="sleeptime"
27             type="Integer"
28             description="time to sleep (in seconds)"
29             value="0"
30             visible="true" />
31     </jobargs>
32
33     <job>
34         <fact name="description"
35             type="String"
36             value="This example job does absolutely nothing. It just
returns immediately. For testing network latency." />
37     </job>
38
39 </policy>
40

```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Job Details

The Quickie job can be broken down into the following separate operations:

- ♦ “zosadmin deploy” on page 124
- ♦ “job_started_event” on page 124
- ♦ “joblet_started_event” on page 124

zosadmin deploy

The job is first deployed into the grid, as shown in lines 20-29 of [quickie.jdl \(page 121\)](#). When jobs are deployed into the grid, they can optionally be organized for grouping. In this example, the Quickie job is added to the group named `examples` and displays in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples
```

For a general overview of how jobs are added to groups during deployment, see “[Walkthrough: Deploying a Sample Job](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

job_started_event

As shown in line 37 of [quickie.jdl \(page 121\)](#), scheduling one or more instances of the Quickie joblet to run immediately is the second operation performed by the Quickie job. When the Quickie job class receives a `job_started_event()` notification, it schedules the number of QuickieJoblet instances as indicated by the value of the setting `numJoblets`, whose value might have been supplied on the command line or from the `quickie.policy` file (see line 20-24 in [quickie.policy \(page 122\)](#)).

joblet_started_event

The final operation performed by the Quickie job is for the joblet to sleep an amount of time as specified by the value of the setting `sleeptime` (see line 48 in [quickie.jdl \(page 121\)](#)), and then exit.

Configure and Run

1 Deploy `quickie.job` into the grid:

```
> zosadmin deploy quickie.job
```

2 Display the list of deployed jobs:

```
> zos joblist
```

quickie should appear in this list.

3 Run the job on one or more resources using the default values for `numJoblets` and `sleeptime`:

```
> zos run quickie
```

4 Run the job on one or more resources using supplied values for `numJoblets` and `sleeptime`:

```
> zos run quickie numJoblets=10 sleeptime=3
JobID: zenuser.quickie.418
```

```
> zos status zenuser.quickie.418
Completed
```

```
> zos log zenuser.quickie.418
Launching 10 joblets
```

Ten joblets will be run simultaneously, depending on the number of resources available in the grid and how many simultaneous jobs each resource is configured to run. After the job runs, each quickie joblet instance simply starts up, sleeps for 3 seconds, and then exits.

See Also

- ◆ Setting Constraints Using Policies (see [Section 4.4, “Policy Management,” on page 34](#) and [Chapter 5, “Policy Elements,” on page 45](#)).
- ◆ Adding jobs to groups during deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job \(page 168\)](#)).
- ◆ Scheduling multiple instances of a joblet

10.5 General Purpose Jobs

The following examples demonstrate general purpose job concepts:

- ◆ [“dgtest.job” on page 126](#)
- ◆ [“failover.job” on page 136](#)
- ◆ [“instclients.job” on page 143](#)
- ◆ [“notepad.job” on page 150](#)
- ◆ [“sweeper.job” on page 155](#)
- ◆ [“whoami.job” on page 162](#)

dgtest.job

This job demonstrates downloading a file from the datagrid.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
Logged into grid as zenuser

> zos jobinfo --detail dgtest
Jobname/Parameters      Attributes
-----
dgtest                  Desc: This job demonstrates downloading from the Datagrid

      multicast         Desc: Whether to download using multicast or unicast
                        Type: Boolean
                        Default: false

      filename          Desc: The filename to download from the Datagrid
                        Type: String
                        Default: None! Value must be specified
```

Description

Demonstrates usage of the datagrid to download a file stored on the PlateSpin Orchestrate Server to a node. For additional background information, see [Section 3.1, “Defining the Datagrid,” on page 25](#).

Because it typically grows quite large, the physical location of the PlateSpin Orchestrate root directory is important. Use the following procedure to determine the location of the datagrid in the Orchestrate Development Client:

- 1 Select the grid id on the left in the PlateSpin Orchestrate Explorer window >
- 2 Click the *Constraints/Facts* tab.

The read-only fact name (`matrix.datagrid.root`) is located here by default:

```
/var/opt/novell/zenworks/zos/server
```

The top level directory name is `dataGrid`.

Contents of the PlateSpin Orchestrate can be seen with the command:

```
> zos dir grid:///
<DIR>      Dec-6-2007 6:55 installs
<DIR>      Dec-6-2007 6:55 jobs
<DIR>      Dec-6-2007 22:01 users
<DIR>      Dec-6-2007 6:55 vms
<DIR>      Dec-6-2007 6:56 warehouse
```

Job Files

The files that make up the Dgtest job include:

```

dgtest                               # Total: 238 lines
|-- dgtest.jdl                       # 172 lines
`-- dgtest.policy                    # 66 lines

```

dgtest.jdl

```

1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY ,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: dgtest.jdl,v 1.4 2008/03/05 20:05:43 ray Exp $
13 # -----
-----
14
15 """
16 Example usage of DataGrid to download a file stored on the Server to a
node.
17
18 Setup:
19     Before running the job, you must:
20         (1) Create a dgtest resource group using the management console.
21         (2) Copy a suitable file into the Server DataGrid
22         (3) Modify the dgtest policy with the filename to download
23             (to not use the default test file).
24
25     For example, use the following command to copy the file 'suse-10-fla
t.vmdk'
26         into the deployment area for the job 'dgtest'
27         >zos mkdir grid:///images
28
29         >zos copy suse-10-flat.vmdk grid:///images/
30
31     To verify the file is there:
32         >zos dir grid:///images
33
34
35     To start the job after the above setup steps are complete:
36         >zos run dgtest filename=suse-10-flat.vmdk
37
38 """
39 import os,time
40
41 #
42 # Add to the 'examples' group on deployment
43 #

```

```

44 if __mode__ == "deploy":
45     try:
46         jobgroupname = "examples"
47         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
48         if jobgroup == None:
49             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
50             jobgroup.addMember(__jobname__)
51     except:
52         exc_type, exc_value, exc_traceback = sys.exc_info()
53         print "Error adding %s to %s group: %s %s" % (__jobname__, jobgr
oupname, exc_type, exc_value)
54
55
56 class test(Job):
57
58     def job_started_event(self):
59         filename = self.getFact("jobargs.filename")
60         print "Starting Datagrid Test Job."
61         print "Filename: %s" % (filename)
62
63         rg = None
64         try:
65             rg = getMatrix().getGroup("resource", "dgtest")
66         except:
67             # no such group
68             pass
69
70         if rg == None:
71             self.fail("The resource group 'dgtest' was not found. It is
required for this job.")
72             return
73
74         members = rg.getMembers()
75         count = 0
76         for resource in members:
77             if resource.getFact("resource.online") == True and \
78                 resource.getFact("resource.enabled") == True:
79                 count += 1
80
81         memo = "Scheduling Datagrid Test on %d Joblets" % (count)
82         self.setFact("jobinstance.memo", memo)
83         print memo
84         self.schedule(testnode, count)
85
86
87 class testnode(Joblet):
88
89     def joblet_started_event(self):
90         jobletnum = self.getFact("joblet.number")
91         print "Running datagrid test joblet #%d" % (jobletnum)
92         filename = self.getFact("jobargs.filename")
93         multicast = self.getFact("jobargs.multicast")
94
95         # Test download a file from server job directory
96         dg_url = "grid:///images/" + filename
97
98         # Create an instance of the JDL DataGrid object
99         # This object is used to manage DataGrid operations
100        dg = DataGrid()

```



```

101
102     # Set to always force a download.
103     dg.setCache(False)
104
105     # Set whether to use multicast or unicast
106     # If set to True, then the following 4 multicast
107     # options are applicable
108     dg.setMulticast(multicast)
109
110     # how long to wait for a quorum (milliseconds)
111     #dg.setMulticastWait( 10000 )
112
113     # Number of receivers that constitute a quorum
114     #dg.setMulticastQuorum(4)
115
116     # Requested data rate in bytes per second. 0 means use default
117     #dg.setMulticastRate(0)
118
119     # Min number of receivers
120     #dg.setMulticastMin(1)
121
122     if multicast:
123         mode = "multicast"
124     else:
125         mode = "unicast"
126
127     memo = "Starting %s download of file: %s" % (mode,dg_url)
128     self.setFact("joblet.memo",memo)
129     print memo
130
131     # Destination defaults to Node's Joblet dir.
132     # Change this path to go to any other local filesystem.
133     # e.g. to store in /tmp:
134     #     dest = "/tmp/" + filename
135     dest = filename
136     try:
137         dg.copy(dg_url,dest)
138     except:
139         exc_type, exc_value, exc_traceback = sys.exc_info()
140         retryUnicast = False
141         if multicast == True:
142             # If node's OS and/or NIC does not fully support multi
cast,
143             # then the node will timeout waiting for broadcasts.
144             # Note the error and fallback to unicast
145             if exc_type != None and len(str(exc_type)) > 0:
146                 msg = str(exc_type)
147                 index = msg.find("Multicast receive timed out")
148                 retryUnicast = index != -1
149
150         if retryUnicast:
151             memo = "Multicast timeout. Fallback to unicast"
152             self.setFact("joblet.memo",memo)
153             print memo
154             dg.setMulticast(False)
155             dg.copy(dg_url,dest)
156         else:
157             raise exc_type,exc_value
158

```

```

159         if os.path.exists(dest):
160             print dg_url + " downloaded successfully."
161
162             # Show directory listing of downloaded file to job log
163             if self.getFact("resource.os.family") == "windows":
164                 cmd = "dir %s" % (dest)
165             else:
166                 cmd = "ls -lsart %s" % (dest)
167
168             system(cmd)
169         else:
170             raise RuntimeError, "Datagrid copy() failed"
171
172     print "Datagrid test completed"

```

dgtest.policy

```

1  <!--
2
=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT.  NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
=====
=
13 * $Id: dgtest.policy,v 1.2 2008/02/27 20:49:29 john Exp $
14
=====
=
15 -->
16
17 <policy>
18
19     <jobargs>
20
21         <!--
22             Name of file that is stored in the Datagrid area to
23             download to the resource.
24
25             A value for this fact the 'zos run' is assigned when
26             using the 'zos run' command.
27         -->
28         <fact name="filename"
29             type="String"
30             description="The filename to download from the Datagrid"
31         />

```

```

32
33     <fact name="multicast"
34         type="Boolean"
35         description="Whether to download using multicast or unicast"
36         value="false" />
37
38 </jobargs>
39
40 <job>
41     <fact name="description"
42         type="String"
43         value="This job demonstrates downloading from the Datagrid"
44     />
45     <!-- limit to one per host -->
46     <fact name="joblet.maxperresource"
47         type="Integer"
48         value="1" />
49 </job>
50
51
52 <!--
53     This job will only run on resources in the "dgtest" resource group.
54
55     You must create a Resource Group named 'dgtest' using the
management
56     console and populate the new group with resources that you wish to
have
57     participate in the datagrid test.
58     -->
59     <constraint type="resource" reason="No resources are in the dgtest
group" >
60
61         <contains fact="resource.groups" value="dgtest"
62             reason="Resource is not in the dgtest group" />
63
64     </constraint>
65
66 </policy>

```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

test

Class test (line 56 in [dgtest.jdl \(page 127\)](#) is derived from the [Job](#) class.

testnode

Class testnode (line 87 in [dgtest.jdl \(page 127\)](#) is derived from the [Joblet \(page 228\)](#) class.

Job Details

dgtest.job can be broken down into the following parts:

- ♦ [“Policy” on page 132](#)
- ♦ [“zosadmin deploy” on page 132](#)
- ♦ [“job_started_event” on page 133](#)
- ♦ [“joblet_started_event” on page 133](#)

Policy

In addition to describing the `filename` and `multicast` jobargs and the default settings for `multicast` (lines 19-38) in the [dgtest.policy \(page 130\)](#) file, there is the `<job/>` section (lines 40-49), which describes static facts ([Section 5.2, “Facts,” on page 45](#)). You must assign the `filename` argument when executing this example. This is only the name of the file in the “images” area of PlateSpin Orchestrate. For example, for `grid:///images/disk.img`, just assign `disk.img` to the argument. This file must be in the PlateSpin Orchestrate file system for fetching and delivering to remote nodes used in this example.

To populate the PlateSpin Orchestrate, use the `zos copy` command. For example, for a file named `suse-9-flat.vmd` in the current directory, use the following command:

```
> zos mkdir grid:///images
> zos copy suse-9-flat.vmd grid:///images/
```

The `multicast` jobarg is a Boolean, defaulted to `false` so that unicast is used for transport. Set this value to `true` to use multicast transport for delivery of the file.

The policy also describes a `resource.groups` constraint. (For more information, see [“Constraints” on page 45](#)). This requires a resource group named `dgtest` (lines 52-64 in [dgtest.policy \(page 130\)](#)) and that group should have member nodes. Consequently, you must create this resource group using the Orchestrate Server Development Client and assign it some nodes to run this example successfully.

zosadmin deploy

When the Orchestrate Server deploys a job for the first time (see [Section 7.5, “Deploying Jobs,” on page 53](#)), the job JDL files are executed in a special deploy mode. Looking at [dgtest.jdl \(page 127\)](#), you might notice that when the job is deployed (line 44), either through the Orchestrate

Development Client or the `zosadmin deploy` command, that it attempts to find the `examples` jobgroup (lines 46-47), create it if missing (lines 48-49), and add the `dgtest` job to the group (line 50).

If this deployment fails for some reason, an exception is thrown (line 51), which prints (line 53) the job name, group name, exception type, and value.

job_started_event

In `dgtest.jdl` (page 127), the `test` class (line 56) defines only the required `job_started_event` (line 58) method. This method runs on the Orchestration server when the job is run to launch the joblets.

When `job_started_event` is executed, it gets the name of the file assigned to the `jobargs.filename` variable and prints useful tracing information (lines 59-61). It then tries to find the resource group named `dgtest`. If the resource group does not exist, the member `fail` string is set to inform the user and returns without scheduling the joblet(s) (lines 63-72).

After finding the `dgtest` group, the job gets the member list and determines how many nodes are online and enabled. The total `count` is stored in lines 74-79. After setting the memo line in the Orchestrate Development Client (81-82), the job schedules `count` number of `testnode` joblets (line 84).

joblet_started_event

In `dgtest.jdl` (page 127), the `testnode` class (line 87) defines only the required `joblet_started_event` (line 89) method. This method runs on the Orchestrate Agent nodes when scheduled by a `Job` (page 226) class.

The `joblet_started_event` prints some trace information (lines 90-91), gets the name of the file to transfer (line 92) and the mode of transfer (line 93), and creates the grid URL for the file (line 96).

A `DataGrid` (page 216) is instantiated (line 100), set not to cache (line 103), and set to use the multicast jobarg (line 108). The next four settings control multicast behavior are commented out (lines 111, 114, 117, and 120).

The joblet prints a memo line for the Orchestrate Development Client (lines 122-128), sets the location for the file on the local node (line 135), and tries to transfer the file from the datagrid (line 137).

If the datagrid copy at line 137 fails for some reason, we have a retry mechanism in the exception handler (lines 138-157). The information for why the exception occurred is fetched (line 139).

The variable `retryUnicast` (line 140) is set `False` and will only be set `True` if the failed download attempt was using multicast transport and the exception type has the string "Multicast receive timed out" (lines 140-148). If the `timed out` string is not found, the triad assigns the `retryUnicast` a value of `-1`. With this logic, either multicast timeout or not, a unicast attempt is made if multicast fails.

If you get to line 150 from a failed multicast copy, a memo for the Orchestrate Development Client is set and printed to the log (151-152), `setMulticast` is set to `false` (154), and another copy from the datagrid is attempted.

If we get to line 150 from a failed unicast copy, an exception is raised (line 157) and we're done.

Configure and Run

```
> zos run dgtest filename=suse-9-flat.vmd
JobID: zenuser.dgtest.323
```

Looks like it ran successfully; let's see what the log says:

```
> zos log zenuser.dgtest.323
Starting Datagrid Test Job.
Filename: suse-9-flat.vmd
Job 'zenuser.dgtest.323' terminated because of failure. Reason: The resource
group 'dgtest' was not found. It is required for this job.
```

There is no resource group. Using the Orchestration Development Client create the resource group dgtest:

```
> zos run dgtest filename=suse-9-flat.vmd
JobID: zenuser.dgtest.324
```

```
> zos log zenuser.dgtest.324
Starting Datagrid Test Job.
Filename: suse-9-flat.vmd
Scheduling Datagrid Test on 0 Joblets
```

NOTE: No joblets were scheduled because we have no active nodes in the group.

Using the Orchestration Development Client, populate the dgtest group with nodes that are both online and enabled:

```
> zos run dgtest filename=suse-9-flat.vmd
JobID: zenuser.dgtest.325
```

```
> zos log zenuser.dgtest.325
Starting Datagrid Test Job.
Filename: suse-9-flat.vmd
Scheduling Datagrid Test on 2 Joblets
[freeze] Running datagrid test joblet #0
[freeze] Starting unicast download of file: grid:///images/suse-9-flat.vmd
[freeze] Traceback (innermost last):
[freeze]   File "dgtest.jdl", line 170, in joblet_started_event
[freeze] copy() failed: DataGrid file "/images/suse-9-flat.vmd" does not
exist.
Job 'zenuser.dgtest.325' terminated because of failure. Reason: Job failed
because of too many joblet failures (job.joblet.maxfailures = 0)
[melt] Running datagrid test joblet #1
[melt] Starting unicast download of file: grid:///images/suse-9-flat.vmd
[melt] Traceback (innermost last):
[melt]   File "dgtest.jdl", line 170, in joblet_started_event
[melt] copy() failed: DataGrid file "/images/suse-9-flat.vmd" does not exist.
```

Because the path and the file in the DataGrid are missing, we need to create and populate them:

```
> zos mkdir grid:///images
Directory created.

> zos copy suse-9-flat.vmd grid:///images/
suse-9-flat.vmd copied.
```

```
> zos run dgtest filename=suse-9-flat.vmd
JobID: zenuser.dgtest.326

> zos log zenuser.dgtest.326
Starting Datagrid Test Job.
Filename: suse-9-flat.vmd
Scheduling Datagrid Test on 2 Joblets
[melt] Running datagrid test joblet #1
[melt] Starting unicast download of file: grid:///images/suse-9-flat.vmd
[melt] grid:///images/suse-9-flat.vmd downloaded successfully.
[melt] 16732 -rw-r--r-- 1 root root 17108462 Dec 21 21:32 suse-9-flat.vmd
[melt] Datagrid test completed
[freeze] Running datagrid test joblet #0
[freeze] Starting unicast download of file: grid:///images/suse-9-flat.vmd
[freeze] grid:///images/suse-9-flat.vmd downloaded successfully.
[freeze] 16732 -rw-r--r-- 1 root root 17108462 Dec 21 21:31 suse-9-flat.vmd
[freeze] Datagrid test completed
```

Finally, the file is deployed from the datagrid and copied successfully. However, you will not find it if you look for it on the agent after the joblet is finished. By default, the file is deployed only for the joblet's lifetime into a directory for the joblet, like the following:

```
/var/opt/novell/zenworks/zos/agent/node.default/melt/zenuser.dgtest.326.0
```

So, for a more permanent demonstration, see lines 132-134 in [dgtest.jdl \(page 127\)](#). Uncomment line 134 and comment out line 135 to store your file in the `/tmp` directory and have it continue to exist on the agent after the joblet executes completely.

failover.job

A test job that demonstrates handling of joblet failover.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
  Logged into grid as zenuser

> zos jobinfo --detail failover
Jobname/Parameters      Attributes
-----
failover                Desc: This test jobs can be used to demonstrate joblet
                        failover handling.

      sleeptime         Desc: specify the execute length of joblet before failure in
                        seconds
                        Type: Integer
                        Default: 7

      numJoblets        Desc: joblets to run
                        Type: Integer
                        Default: 1
```

Description

Schedules one joblet, which fails, then re-instantiates in a repeating cycle until a specified retry limit is reached and the Orchestration Server does not create another instance. This example demonstrates how the orchestration server can be made more robust, as described in [Section 7.13, “Improving Job and Joblet Robustness,” on page 65](#).

The files that make up the Failover job include:

```
failover                # Total: 94 lines
|-- failover.jdl        #   64 lines
`-- failover.policy     #   30 lines
```

failover.jdl

```
1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
```



```

-----
12 # $Id: failover.jdl,v 1.3 2008/02/27 20:50:00 john Exp $
13 # -----
-----
14
15 # Test job to illustrate joblet failover and max retry limits
16 #
17 # Job args:
18 #   numJoblets - specify number of Joblets to run
19 #   sleeptime -- specify the execute length of joblet before failure in
seconds
20 #
21
22 import sys,os,time
23
24 #
25 # Add to the 'examples' group on deployment
26 #
27 if __mode__ == "deploy":
28     try:
29         jobgroupname = "examples"
30         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
31         if jobgroup == None:
32             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
33         jobgroup.addMember(__jobname__)
34     except:
35         exc_type, exc_value, exc_traceback = sys.exc_info()
36         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
37
38
39 class failover(Job):
40
41     def job_started_event(self):
42         numJoblets = self.getFact("jobargs.numJoblets")
43         print 'Launching ', numJoblets, ' joblets'
44         self.schedule(failoverjoblet,numJoblets)
45
46
47 class failoverjoblet(Joblet):
48
49     def joblet_started_event(self):
50         print "----- joblet_started_event"
51         print "node=%s joblet=%d" % (self.getFact("resource.id"),
self.getFact("joblet.number"))
52         print "self.getFact(joblet.retrynumber)=%d" %
(self.getFact("joblet.retrynumber"))
53         print "self.getFact(job.joblet.maxretry)=%d" %
(self.getFact("job.joblet.maxretry"))
54
55         sleeptime = self.getFact("jobargs.sleeptime")
56         print "sleeping for %d seconds" % (sleeptime)
57         time.sleep(sleeptime)
58
59         # This will cause joblet failure and thus retry
60         raise RuntimeError, "Artificial error in joblet. node=%s" %

```

```
(self.getFact("resource.id"))
61
62
63
64
```

failover.policy

```
1  <!--
2
*=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT. NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
*=====
=
13 * $Id: failover.policy,v 1.2 2008/02/27 20:50:00 john Exp $
14
*=====
=
15 -->
16
17 <policy>
18   <jobargs>
19     <fact name="sleeptime" description="specify the execute length of
joblet before failure in seconds" value="7" type="Integer" />
20     <fact name="numJoblets" description="joblets to run" value="1"
type="Integer" />
21   </jobargs>
22
23   <job>
24     <fact name="description" value="This test jobs can be used to
demonstrate joblet failover handling." type="String" />
25
26     <!-- Number of times to retry joblet on failure -->
27     <fact name="joblet.maxretry" type="Integer" value="3" />
28   </job>
29 </policy>
30
```

Classes and Methods

Definitions:

Class `failover` in line 25 of `failover.jdl` (page 136) is derived from the `Job` (page 226) class; and the class `failoverjoblet` in line 33 of `failover.jdl` (page 136) is derived from the `Joblet` (page 228) class.

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. `MatrixInfo` is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

failover

Class `failover` (line 39 in `dgtest.jdl` (page 127)) is derived from the `Job` class.

failoverjoblet

Class `failoverjoblet` (line 47 in `dgtest.jdl` (page 127)) is derived from the `Joblet` (page 228) class.

Job Details

The following sections describe the Failover job:

- ♦ “zosadmin deploy” on page 139
- ♦ “job_started Event” on page 140
- ♦ “job_started Event” on page 140

zosadmin deploy

In `failover.policy` (page 138), in addition to describing the jobargs and default settings for `sleeptime` and `numJoblets` (lines 18-21), the `<job/>` section (lines 23-28) describes static facts (see Section 5.2, “Facts,” on page 45). Note that the `joblet.maxretries` attribute in line 27 has a default setting of 0 but is set here to 3. This attribute can also be modified in the `failover.jdl` (page 136) file by inserting a line between line 41 and 42, as shown in the following example:

```
41     def job_started_event(self):
++         self.setFact("job.joblet.maxretries", 3)
42         numJoblets = self.getFact("jobargs.numJoblets")
```

job_started Event

After the Orchestrate Server deploys a job for the first time (see [Section 7.5, “Deploying Jobs,” on page 53](#)), the job JDL files are executed in a special “deploy” mode. When the job is deployed (line 27, [failover.jdl \(page 136\)](#), it attempts to find the `examples` jobgroup (lines 29-30), creates it if it is missing (lines 31-32), and adds the failover job to the group (line 33).

Jobs can be deployed using either the Orchestrate Development Client or the `zosadmin deploy` command. If the deployment fails for some reason, an exception is thrown (line 34), which prints the job name (line 36), group name, exception type, and value.

job_started Event

In [failover.jdl \(page 136\)](#), the failover class (line 39) defines only the required `job_started_event` (line 41) method. This method runs on the Orchestrate Server when the job is run to launch the joblets.

On execution, the `job_started_event` simply gets the number of joblets to create (`numJoblets` in line 42), then schedules that specified number of instances (line 44) of the `failoverjoblet` class. The `failoverjoblet` class (lines 47-60) defines only the required `joblet_started_event` (line 49) method.

When executed on an agent node, the `joblet_started_event` prints some helpful information for tracking execution (lines 50-53). The first output is where the joblet is running and which instance is running (line 51). The current joblet retry number (line 52) is displayed, followed by the job’s static `joblet.maxretry` (line 53) that was specified in the policy file.

The joblet then sleeps for `jobargs.sleeptime` seconds (lines 55-57) and on waking raises an exception of type `RuntimeError` (line 60).

This is the point of this example. After a `RuntimeError` exception is thrown, the zos server attempts to run the same instance of the joblet again if `job.joblet.maxretry` (default is 0) is less than or equal to `joblet.retrynumber`.

Configure and Run

You must be logged into the Orchestrate Server before you run `zosadmin` or `zos` commands.

1 Deploy failover.job into the grid:

```
> zosadmin deploy failover.job
JobID: zenuser.failover.269
```

The job appears to have run successfully, now take a look at the log and see the joblet failure and being relaunched until finally the “maxretry” count is exceeded and the job exits with a failure status:

2 Display the list of deployed jobs:

```
> zos joblist
```

failover should appear in this list.

3 Run the job on one or more resources using the default values for numJoblets and sleeptime, specified in the failover.policy file:

```
> zos run failover sleeptime=1 numJoblets=2
JobID: zenuser.failover.269
```

The job appears to have run successfully, now take a look at the log and see the joblet failure and being relaunched until finally the maxretry count is exceeded and the job exits with a failure status:

```
> zos log zenuser.failover.269Launching 2 joblets
[melt] ----- joblet_started_event
[melt] node=melt joblet=1
[melt] self.getFact(joblet.retrynumber)=0
[melt] self.getFact(job.joblet.maxretry)=3
[melt] sleeping for 1 seconds
[melt] Traceback (innermost last):
[melt]   File "failover.jdl", line 60, in joblet_started_event
[melt] RuntimeError: Artifical error in joblet. node=melt
[freeze] ----- joblet_started_event
[freeze] node=freeze joblet=0
[freeze] self.getFact(joblet.retrynumber)=0
[freeze] self.getFact(job.joblet.maxretry)=3
[freeze] sleeping for 1 seconds
[freeze] Traceback (innermost last):
[freeze]   File "failover.jdl", line 60, in joblet_started_event
[freeze] RuntimeError: Artifical error in joblet. node=freeze
[melt] ----- joblet_started_event
[melt] node=melt joblet=0
[melt] self.getFact(joblet.retrynumber)=1
[melt] self.getFact(job.joblet.maxretry)=3
[melt] sleeping for 1 seconds
[melt] Traceback (innermost last):
[melt]   File "failover.jdl", line 60, in joblet_started_event
[melt] RuntimeError: Artifical error in joblet. node=melt
[freeze] ----- joblet_started_event
[freeze] node=freeze joblet=1
[freeze] self.getFact(joblet.retrynumber)=1
[freeze] self.getFact(job.joblet.maxretry)=3
[freeze] sleeping for 1 seconds
[freeze] Traceback (innermost last):
[freeze]   File "failover.jdl", line 60, in joblet_started_event
[freeze] RuntimeError: Artifical error in joblet. node=freeze
[melt] ----- joblet_started_event
[melt] node=melt joblet=1
[melt] self.getFact(joblet.retrynumber)=2
[melt] self.getFact(job.joblet.maxretry)=3
[melt] sleeping for 1 seconds
[melt] Traceback (innermost last):
[melt]   File "failover.jdl", line 60, in joblet_started_event
[melt] RuntimeError: Artifical error in joblet. node=melt
[freeze] ----- joblet_started_event
[freeze] node=freeze joblet=0
[freeze] self.getFact(joblet.retrynumber)=2
[freeze] self.getFact(job.joblet.maxretry)=3
[freeze] sleeping for 1 seconds
[freeze] Traceback (innermost last):
[freeze]   File "failover.jdl", line 60, in joblet_started_event
[freeze] RuntimeError: Artifical error in joblet. node=freeze
```

See Also

- ◆ [Setting Constraints Using Policies \(Section 4.4, “Policy Management,” on page 34 and Chapter 5, “Policy Elements,” on page 45\).](#)

- ◆ Adding Jobs to Groups During Deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job \(page 168\)](#)).
- ◆ Executing Commands Using [Exec \(page 219\)](#)

instclients.job

Installs the PlateSpin Orchestrate client applications to the specified resource machine. Note that while most of the other examples are deployed by default, this example is not.

Detail

The following concepts are demonstrated:

- ♦ Using constraints to restrict joblet execution to a specific resource.
- ♦ Adding files to a job's directory in the datagrid, and retrieving them during joblet execution.
- ♦ Using the **Exec** class to send a command to the operating system. The system command is invoked directly without using the system command interpreter (either `cmd.exe` or `/bin/sh`).

Usage

```
> zosadmin login --user zosadmin Login to server: skate
Please enter current password for 'zosadmin':
Logged into grid on server 'skate'

> cd /opt/novell/zenworks/zos/server/examples
> zosadmin deploy instclients.job
instclients successfully deployed

> zos login --user zenuser
Please enter current password for 'zenuser':
  Logged into grid as zenuser

> zos jobinfo --detail instclients
Jobname/Parameters  Attributes
-----
instclients        Desc: This job installs the ZOS clients on a resource

      host          Desc: The host name of resource to install on
                    Type: String
                    Default: None! Value must be specified
```

Description

The files that make up the instclients job include:

```
instclients          # Total: 138 lines
|-- instclients.jdl  #    97 lines
`-- instclients.policy #    41 lines
```

instclients.jdl

```
1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
```

```

AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: instclients.jdl,v 1.4 2008/03/05 20:05:54 ray Exp $
13 # -----
-----
14
15 """
16
17 Run install clients on a resource
18
19 Setup:
20     Before running the job, you must copy installers into DataGrid of
21     server.
22
23     >zos copy zosclients_windows_1_3_0_with_jre.exe grid:///!\instclients/
24
25 """
26 import os,time
27
28 #
29 # Add to the 'examples' group on deployment
30 #
31 if __mode__ == "deploy":
32     try:
33         jobgroupname = "examples"
34         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
35         if jobgroup == None:
36             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
37         jobgroup.addMember(__jobname__)
38     except:
39         exc_type, exc_value, exc_traceback = sys.exc_info()
40         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
41
42
43 class InstClients(Job):
44
45     def job_started_event(self):
46         print "Scheduling joblet"
47         self.schedule(InstClientsJoblet)
48
49
50 class InstClientsJoblet(Joblet):
51
52     def joblet_started_event(self):
53         print "Launching Installer"
54         windowsInstaller = "zosclients_windows_1_3_0_with_jre.exe"
55         linuxInstaller = "zosclients_linux_1_3_0_with_jre.sh"
56         if self.getFact("resource.os.family") == "windows":
57             print "Downloading Windows install"
58             dg = DataGrid()

```



```

59         dg.copy("grid:///!instclients/" +
windowsInstaller, windowsInstaller)
60
61         print "Starting install"
62         cmd = self.getcwd() + "/" + windowsInstaller + " -q "
63         e = Exec()
64         e.setCommand(cmd)
65         e.setRunAsJobUser(False)
66         e.writeStdoutToLog()
67         e.writeStderrToLog()
68         result = e.execute()
69     else:
70         print "Downloading Linux install"
71         dg = DataGrid()
72         dg.copy("grid:///!instclients/" +
linuxInstaller, linuxInstaller)
73
74         print "Starting install"
75         cmd = "chmod +x " + self.getcwd() + "/" + linuxInstaller
76         print "cmd=%s" % (cmd)
77         e = Exec()
78         e.setCommand(cmd)
79         e.setRunAsJobUser(False)
80         e.writeStdoutToLog()
81         e.writeStderrToLog()
82         result = e.execute()
83
84         cmd = self.getcwd() + "/" + linuxInstaller + " -q"
85         print "cmd=%s" % (cmd)
86         e = Exec()
87         e.setRunAsJobUser(False)
88         e.setCommand(cmd)
89         e.writeStdoutToLog()
90         e.writeStderrToLog()
91         result = e.execute()
92
93     if result == 0:
94         print "Install complete"
95     else:
96         print "result=%d" % (result)
97

```

instclients.policy

```

1  <!--
2
*=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY ,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT. NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,

```

```

10  * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11  * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
=====
=
13  * $Id: instclients.policy,v 1.2 2008/02/27 20:50:26 john Exp $
14
=====
=
15  -->
16
17  <policy>
18
19      <jobargs>
20
21          <fact name="host"
22              type="String"
23              description="The host name of resource to install on"
24              />
25
26      </jobargs>
27
28      <job>
29          <fact name="description"
30              type="String"
31              value="This job installs the ZOS clients on a resource" />
32      </job>
33
34      <constraint type="resource" >
35
36          <eq fact="resource.id" factvalue="jobargs.host" />
37
38      </constraint>
39
40  </policy>
41

```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Exec

Used to manage command line execution on resources.

DataGrid

Provides a way to interact with the datagrid. Operations include copying files from the datagrid down to the resource for joblet usage and uploading files from a resource to the datagrid.

Job Details

The following sections describe the instclients job:

- ♦ [“zosadmin deploy” on page 147](#)
- ♦ [“job_started_event” on page 147](#)
- ♦ [“joblet_started_event” on page 147](#)

zosadmin deploy

When jobs are deployed into the grid, they can optionally be placed in groups for organization and easy reference. In this case, the instclients job will be added to the group named Examples (lines 31-40), and will show up in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples.
```

For a general overview of how jobs are added to groups during deployment, see [“Walkthrough: Deploying a Sample Job”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

job_started_event

When the instclients job receives a `job_started_event`, it schedules a single instance of the Instclients joblet to be run (see line 45 of [instclients.jdl \(page 143\)](#)). The resource that runs the joblet is determined from the resource constraint specified in [instclients.policy \(page 145\)](#), lines 21-24, and from the value for the parameter `host` supplied on the command line.

joblet_started_event

After the Instclients joblet is executed on a particular resource, it receives a `joblet_started_event`. When this happens, the Instclients joblet decides which Orchestrate Client installation file to download, and the commands to execute on the operating system by checking the value of `resource.os.family` (see line 56 of [instclients.jdl \(page 143\)](#)). The `resource.os.family` fact does not exist in the `instclients.policy` file, but is instead provided by the PlateSpin Orchestrate system.

After deciding which operating system the joblet is being run on, the `Instclients` joblet uses the `DataGrid` class to download the appropriate client installation file to the current working directory of the running joblet (see lines 58-59 and 71-72 in [instclients.jdl \(page 143\)](#)). The URL `grid://!instclients/` points to a directory reserved for the joblet in the datagrid on the server.

After the client installation file has been downloaded from the server, the `Instclients` joblet uses the `Exec` class to begin the installation (see lines 63-68 and 86-91 in `instclients.jdl` (page 143)). As indicated by lines 66, 67, 80, 81, 89 and 90, all standard out and standard err are written to the job's log file.

To view the log file for the `Instclients` job after it has been run, you can execute the command

```
zos log instclients
```

For more information about using `zos`, see [Section 7.5.2, “Using the zosadmin Command Line Tool,” on page 54](#). See the `Exec` class in [PlateSpin Orchestrate Job Classes and JDL Syntax](#) for more information on running commands.

NOTE: The `Instclients` job uses the `Exec` class twice when running on a Linux resource. The first command changes the mode of the installation file to be an executable, and the second runs the installation file.

Configure and Run

Execute the following commands to deploy and run `instclients.job`:

- 1 Copy client installation files into the directory reserved for the `Instclients` joblet in the datagrid of the Orchestrate Server:

```
zos copy zosclients_linux_1_3_0_with_jre.sh grid:///!\instclients/
```

NOTE: Replace “linux” with windows, linux, solaris, etc. for your given operating system, and replace 1_3_0 with your version of the product.

This command copies the file `zosclients_linux_1_3_0_with_jre.sh` into the datagrid job directory for `instclients`.

For more information about using PlateSpin Orchestrate Development Client to copy files, type `zos copy -help`.

NOTE: Replace windows with *linux*, *solaris*, etc. for your given operating system.

- 2 Deploy `instclients.job` into the grid by entering:

```
zosadmin deploy instclients.job
```

- 3 Display the list of deployed jobs by entering:

```
zos joblist
```

`instclients` should appear in this list.

- 4 Run the job on the resource with the given host:

```
zos run instclients host=my_resource_host
```

Installs the Orchestrate clients onto the resource with the host: `my_resource_host`.

See Also

- ♦ [Setting Constraints Using Policies \(Section 4.4, “Policy Management,” on page 34 and Chapter 5, “Policy Elements,” on page 45\)](#).
- ♦ [Adding Jobs to Groups During Deployment \(see how the JDL code can print the ID of group of jobs in `factJunction.job` \(page 168\)\)](#).

- ◆ Scheduling multiple instances of a joblet
- ◆ Setting default parameter values using policies
- ◆ Configuring constraints in a policy file
- ◆ Naming conventions for policy facts ([Section 3.1.1, “PlateSpin Orchestrate Datagrid Filepaths,” on page 25](#).[Section 3.1.1, “PlateSpin Orchestrate Datagrid Filepaths,” on page 25](#))
- ◆ Facts provided by the PlateSpin Orchestrate system that can be referenced within a JDL file
- ◆ Using the PlateSpin Orchestrate Development Client (“[How Do I Interact with PlateSpin Orchestrate?](#)”)
- ◆ Running commands using the `Exec` class.

notepad.job

Launches the Notepad application on a Windows resource.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
  Logged into grid as zenuser

> zos jobinfo --detail notepad
Jobname/Parameters      Attributes
-----
notepad                  Desc: No description available.
```

Description

The files that make up the Notepad job include:

```
notepad                  # Total: 86 lines
|-- notepad.jdl          #   54 lines
`-- notepad.policy       #   32 lines
```

notepad.jdl

```
1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: notepad.jdl,v 1.3 2008/02/27 20:50:47 john Exp $
13 # -----
-----
14
15 ""
16
17 Run Notepad Application on windows resoure
18
19 ""
20 import os,time
21
22 #
23 # Add to the 'examples' group on deployment
24 #
```

```

25 if __mode__ == "deploy":
26     try:
27         jobgroupname = "examples"
28         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
29         if jobgroup == None:
30             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
31             jobgroup.addMember(__jobname__)
32     except:
33         exc_type, exc_value, exc_traceback = sys.exc_info()
34         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
35
36
37 class Notepad(Job):
38
39     def job_started_event(self):
40         print "Scheduling joblet"
41         self.schedule(NotepadJoblet)
42
43
44 class NotepadJoblet(Joblet):
45
46     def joblet_started_event(self):
47         print "Starting Notepad"
48         cmd = "notepad"
49         e = Exec()
50         e.setCommand(cmd)
51         e.writeStdoutToLog()
52         e.writeStderrToLog()
53         result = e.execute()
54

```

notepad.policy

```

1  <!--
2
*=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT.  NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
*=====
=
13 * $Id: notepad.policy,v 1.2 2008/02/27 20:50:47 john Exp $
14
*=====
=

```

```

15  -->
16
17  <policy>
18
19      <constraint type="accept" >
20
21          <gt fact="jobinstance.matchingresources" value="0" reason="No
Windows's resources are available to run Notepad" />
22
23      </constraint>
24
25      <constraint type="resource" >
26
27          <eq fact="resource.os.family" value="windows" reason="Notepad
only runs on Windows OS" />
28
29      </constraint>
30
31  </policy>
32

```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Exec

Used to manage command line execution on resources.

Job Details

The Notepad job is broken down into three separate operations:

- ◆ “zosadmin deploy” on page 153
- ◆ “job_started_event” on page 153
- ◆ “joblet_started_event” on page 153

zosadmin deploy

In [notepad.jdl \(page 150\)](#), lines 25-34 places the job into the “examples” job group. After jobs are deployed into the grid, they can optionally be placed in groups for organization and easy reference. In this case, the Notepad job is added to the group named Examples and appears in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples
```

For a general overview of how jobs are added to groups during deployment, see “[Walkthrough: Deploying a Sample Job](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

job_started_event

Scheduling the Notepad joblet to run immediately is the second operation performed by the Notepad job in line 41 of [notepad.jdl \(page 150\)](#). When the Notepad job class receives a `job_started_event()` notification, it simply schedules the NotepadJoblet class to be run on any target device that meets the restrictions identified in the `notepad.policy` file.

As specified in lines 21 and 27 of [notepad.policy \(page 151\)](#), there must be at least one Windows machine available in the grid for the Notepad job to run. The `accept` constraint in lines 19-23 prevents the Notepad job from being accepted for running if there are no Windows resources available.

The `resource` constraint in lines 25-29 constrain the Orchestrate Job Scheduler to choose a resource that is running a Windows OS only.

For more information on setting constraints using policies, see [Section 4.4, “Policy Management,” on page 34](#) and [Chapter 5, “Policy Elements,” on page 45](#).

joblet_started_event

As specified in lines 49-53 in [notepad.jdl \(page 150\)](#), the joblet executing a command on the target machine is the last operation performed by the Notepad job.

In this example, after the `joblet_started_event()` method of the NotepadJoblet class gets called, the PlateSpin Orchestrate API class named **Exec** is used to run the command `notepad` on is captured and written to the log file for the Notepad job.

Configure and Run

Execute the following commands to deploy and run `notepad.job`:

- 1 Deploy `notepad.job` into the grid:

```
> zosadmin deploy notepad.job
```
- 2 Display the list of deployed jobs:

```
> zos joblist
```

notepad should appear in this list.
- 3 Run the job on the first available Windows resource.

```
> zos run notepad
```

You should now see the Windows Notepad application appear on the screen of the target Windows system. You will see the following error if there are no Windows resources.

```
No Windows resources available to run Notepad
```

See Also

- ♦ Setting Constraints Using Policies see [Section 4.4, “Policy Management,”](#) on page 34 and [Chapter 5, “Policy Elements,”](#) on page 45.
- ♦ Adding Jobs to Groups During Deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job](#) (page 168)).
- ♦ Executing Commands Using [Exec](#) (page 219)

sweeper.job

This example job illustrates how to schedule a "sweep," which is an ordered, serialized scheduling of the joblets across all matching resources.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
  Logged into grid as zenuser

> zos jobinfo --detail sweeper
Jobname/Parameters  Attributes
-----
sweeper             Desc: This example job illustrates how to schedule a 'sweep'
                    accross all matching resources.

                    sleeptime      Desc: time to sleep (in seconds)
                    Type: Integer
                    Default: 1
```

Options

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Exec

Used to manage command line execution on resources.

sleeptime

Specifies the time in seconds that the job remains dormant before running (default 1).

Description

The files that make up the Sweeper job include:

```
sweeper # Total: 140 lines
|-- sweeper.jdl # 66 lines
`-- sweeper.policy # 74 lines
```

The [ScheduleSpec \(page 246\)](#) utility class is also related to this example.

sweeper.jdl

```
1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: sweeper.jdl,v 1.3 2008/02/27 20:51:24 john Exp $
13 # -----
-----
14
15 import time
16
17 #
18 # Add to the 'examples' group on deployment
19 #
20 if __mode__ == "deploy":
21     try:
22         jobgroupname = "examples"
23         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
24         if jobgroup == None:
25             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
26         jobgroup.addMember(__jobname__)
27     except:
28         exc_type, exc_value, exc_traceback = sys.exc_info()
29         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
30
31
32 class sweeperJob(Job):
33
34     def job_started_event(self):
35         self.setFact("jobinstance.memo", self.getFact("job.description"))
36
37         sp = ScheduleSpec()
38
39         # Optionally a constraint can be specified to further limit
matching
40         # resources from the job's default 'resource' constraint. Could
also
```

```

41         # compose an object Constraint.
42         # For example, uncomment to restrict to resource group 'sweeper'
43         #sp.setConstraint("<contains fact='resource.groups'
value='sweeper' />")
44
45         # Specify the joblet to run on each resource
46         sp.setJobletClass(sweeperJoblet)
47
48         # Specify the sweep across active nodes
49         sp.setUseNodeSet(sp.ACTIVE_NODE_SET)
50
51         # Schedule a sweep (creates preassigned joblets)
52         self.scheduleSweep(sp)
53
54         # Now the ScheduleSpec contains the number of joblets created
55         print 'Launched', sp.getCount(), 'joblets'
56
57
58 class sweeperJoblet(Joblet):
59
60     def joblet_started_event(self):
61         msg = "run on resource %s" % (self.getFact("resource.id"))
62         self.setFact("joblet.memo", msg)
63         print "Sweep", msg
64         sleeptime = self.getFact("jobargs.sleeptime")
65         time.sleep(sleeptime)
66

```

sweeper.policy

```

1  <!--
2
=====
=
3  * Copyright (c) 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT. NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
=====
=
13 * $Id: sweeper.policy,v 1.2 2008/02/27 20:51:24 john Exp $
14
=====
=
15 -->
16
17 <policy>
18

```

```

19     <jobargs>
20         <!--
21         - Defines and sets the length of time the joblet should pretend
22         - it is doing something important
23         -->
24         <fact name="sleeptime"
25             type="Integer"
26             description="time to sleep (in seconds)"
27             value="1"
28             visible="true" />
29     </jobargs>
30
31
32     <job>
33         <!--
34         - Give the job a description for GUI's
35         -->
36         <fact name="description"
37             type="String"
38             value="This example job illustrates how to schedule a 'sweep'
across all matching resources." />
39
40         <!--
41         - This activates a built in throttle to limit the number of
42         - resources this job will run on at a time
43         -->
44         <fact name="maxresources"
45             type="Integer"
46             value="3" />
47
48         <!--
49         - Rank resources from least loaded to the highest loaded. The
50         - idea is to run the joblets on the least loaded node first
51         - and hopefully by the time we get to the higher loaded
machines
52         - their load may have gone down
53         -->
54         <!--
55         <fact name="resources.rankby">
56             <array>
57                 <string>resource.loadaverage/a</string>
58             </array>
59         </fact>
60         -->
61
62         <!--
63         - Alternative ranking that is easier to see:
64         - decending alphabetic of node name
65         -->
66         <fact name="resources.rankby">
67             <array>
68                 <string>resource.id/d</string>
69             </array>
70         </fact>
71     </job>
72
73 </policy>
74

```

Classes and Methods

The class `sweeperJob` (see line 32, [sweeper.jdl \(page 156\)](#)) is derived from the **Job Class**.

The class `sweeperJoblet` (see line 58, [sweeper.jdl \(page 156\)](#)) is derived from the **Joblet Class**.

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

Job Details

The `sweeper.job` can be broken down into four separate parts:

- ♦ “Policy” on page 159
- ♦ “zosadmin deploy” on page 159
- ♦ “job_started_event” on page 159
- ♦ “joblet_started_event” on page 160

Policy

In addition to specifying the jobarg and default settings for `sleepTime` in lines 24-28, [sweeper.policy \(page 157\)](#), there also is the `<job/>` section in lines 32-71, which describes static facts (see “Facts” on page 45).

The `resources.rankby` array has two notable setting in this example:

- ♦ **resource.loadaverage:** This is the first string assignment (lines 55-59), which is commented out, that causes joblets to run on the least loaded nodes first. This is the default value and the default launch order for `scheduleSweep`.
- ♦ **resource.id:** This is the second string assignment (lines 66-70), which is actually used, and assigns the string to the rank by array so that joblets run on nodes in reverse alphabetical order.

zosadmin deploy

When the Orchestrate Server deploys a job for the first time (see [Section 7.5, “Deploying Jobs,” on page 53](#)), the job JDL files are executed in a special deploy mode. When `sweeper.jdl` is run in this way (either via the Development Client or the `zosadmin deploy` command), lines 20-29 are executed. This attempts to locate the examples jobgroup (lines 22-23), creates the group if it is not found (lines 24-25), and adds the `sweeper` job to the group (line 26).

If the deployment fails for any reason, then an exception is thrown (line 27), which prints the job name, group name, exception type and value (line 29).

job_started_event

The `sweeperJob` class (line 32) defines only the required `job_started_event` (line 34) method. This method runs on the Orchestrate Server when the job is run to launch the joblets.

When executed, `job_started_event` displays a message on the memo line of the Job Log tab within the Jobs view in the Orchestrate Development Client (line 35), via `jobinstance.memo` (see [Section 7.12.1, “Creating a Job Memo,” on page 63](#)).

Jumping ahead for a moment, instead of calling `self.schedule()` as most the other examples do to instantiate joblets, `sweeperJob` calls `self.scheduleSweep()` (line 52). `scheduleSweep` requires an instance of [ScheduleSpec \(page 246\)](#), so one is created (line 37).

The `ScheduleSpec` method `setConstraint` can be used to constrain the available resources to a particular group, as shown with a comment (line 43). If this `setConstraint` line is uncommented, joblets will only run on members of the `sweeper resource.group` instead of using the default resource group `all`.

NOTE: The `sweeper` group must already be created and have computing nodes assigned to it (see [“Walkthrough: Creating a Resource Account”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*). This constraint would also be ANDed to any existing constraint, including any aggregated policies.

The `sweeperJoblet` is set to be scheduled (line 52), and `setUseNodeSet(intnodeSet)` is assigned (line 49) the value `sp.ACTIVE_NODE_SET`. So, the joblet set is constructed after applying resource constraints to the active/online resources. This in contrast to the other possible value of `sp.PROVISIONABLE_NODE_SET`, where constraints are applied to all provisionable resources.

joblet_started_event

The `sweeperJoblet` class (lines 58-65) defines only the required `joblet_started_event` (line 60) method. After this method is executed, it displays a message on the memo line of the Joblet tab within the Jobs view in the Orchestrate Development Client (lines 61-62). It also prints a similar log message (line 63), and then just sleeps for `jobargs.sleep_time` seconds (lines 64-65) before completion.

Configure and Run

Execute the following commands to deploy and run `sweeper.job`:

1 Deploy `notepad.job` into the grid:

```
> zosadmin deploy sweeper.job
```

2 Display the list of deployed jobs:

```
> zos joblist
```

sweeper should appear in this list.

3 Run the job on one or more resources using the default values for `numJoblets` and `resource`, specified in the `sweeper.policy` file:

```
> zos run sweeper sleeptime=30
```

```
JobID: zenuser.sweeper.420
```

```
> zos status zenuser.sweeper.420
```

```
Completed
```



```
> zos log zenuser.sweeper.420
Launched 3 joblets
[melt] Sweep run on resource melt
[freeze] Sweep run on resource freeze
[skate] Sweep run on resource skate
```

See Also

- ♦ Setting Constraints Using Policies, see [Section 4.4, “Policy Management,”](#) on page 34 and [Chapter 5, “Policy Elements,”](#) on page 45.

whoami.job

Demonstrates using the **Exec** class to send a command to the operating system's default command interpreter. On Microsoft Windows, this is `cmd.exe`. On POSIX systems, this is `/bin/sh`.

Usage

```
> zos login --user zenuser
Please enter current password for 'zenuser':
  Logged into grid as zenuser

~> zos jobinfo --detail whoami
Jobname/Parameters  Attributes
-----
whoami              Desc: This is a demo example of enhanced exec

      numJoblets    Desc: The number of joblets to schedule
                        Type: Integer
                        Default: 1

      resource      Desc: The resource id to run on
                        Type: String
                        Default: .*
```

Description

The files that make up the Whoami job include:

```
whoami                # Total: 118 lines
|-- whoami.jdl        #   69 lines
`-- whoami.policy     #   49 lines
```

whoami.jdl

```
1 # -----
-----
2 # Copyright (c) 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: whoami.jdl,v 1.3 2008/02/27 20:51:34 john Exp $
13 # -----
-----
14
15 """
```

```

16
17 Demonstrate running setuid exec.
18
19 """
20 import os,time
21
22 #
23 # Add to the 'examples' group on deployment
24 #
25 if __mode__ == "deploy":
26     try:
27         jobgroupname = "examples"
28         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
29         if jobgroup == None:
30             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
31             jobgroup.addMember(__jobname__)
32     except:
33         exc_type, exc_value, exc_traceback = sys.exc_info()
34         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
35
36
37 class Whoami(Job):
38
39     def job_started_event(self):
40         # Launch the joblets
41         numJoblets = self.getFact("jobargs.numJoblets")
42         user = self.getFact("user.id")
43         print "Launching %d joblets for user '%s'" % (numJoblets,user)
44         self.schedule(WhoamiJoblet,numJoblets)
45
46
47 class WhoamiJoblet(Joblet):
48
49     def joblet_started_event(self):
50         if self.getFact("resource.os.family") == "windows":
51             cmd = "echo %USERNAME%"
52         elif self.getFact("resource.os.family") == "solaris":
53             cmd = "echo $USER"
54         else:
55             cmd = "whoami"
56         print "cmd=%s" % (cmd)
57
58         # example using built-in system()
59         #result = system(cmd)
60
61         # example using Exec class
62         e = Exec()
63         e.setShellCommand(cmd)
64         e.writeStdoutToLog()
65         e.writeStderrToLog()
66         result = e.execute()
67
68         print "result=%d" % (result)
69

```

whoami.policy

```
1  <!--
2
=====
=
3  * Copyright (c) 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY ,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT.  NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
=====
=
13 * $Id: whoami.policy,v 1.2 2008/02/27 20:51:34 john Exp $
14
=====
=
15  -->
16
17  <policy>
18
19      <jobargs>
20
21          <fact name="numJoblets"
22              type="Integer"
23              description="The number of joblets to schedule"
24              value="1" />
25
26          <fact name="resource"
27              type="String"
28              description="The resource id to run on"
29              value=".*" />
30
31      </jobargs>
32
33      <job>
34          <fact name="description"
35              type="String"
36              value="This is a demo example of enhanced exec" />
37
38          <!-- only allow one run resource at a time so that multiple re
sources can be visited -->
39          <fact name="joblet.maxperresource"
40              type="Integer"
41              value="1" />
42      </job>
43
44      <constraint type="resource" >
45          <eq fact="resource.id" factvalue="jobargs.resource" match="regex p"
```

```
</>
46     </constraint>
47
48 </policy>
49
```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Exec

Used to manage command line execution on resources.

Job Details

The following sections describe the Whoami job:

- ♦ [“zosadmin deploy” on page 165](#)
- ♦ [“job_started_event” on page 166](#)
- ♦ [“joblet_started_event” on page 166](#)

zosadmin deploy

When jobs are deployed into the grid, they can optionally be placed in groups for organization and easy reference. In this case, the Whoami job is added to the group named “examples” (see lines 25-34 of `whoami.jdl`) and is displayed in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples.
```

For a general overview of how jobs are added to groups during deployment, see [“Walkthrough: Deploying a Sample Job”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

job_started_event

When the Whoami job receives a `job_started_event`, it schedules one or more instances of the Whoami joblet to be run (see line 44 in [whoami.jdl \(page 162\)](#)). The number of WhoamiJoblet instances is indicated by the value of the `numJoblets` fact, whose value might have been supplied on the command-line, or referenced from what's been supplied in the `whoami.policy` file by default (see lines 21-24 in [whoami.policy \(page 164\)](#)).

In addition to supplying a default value for `numJoblets`, the `whoami.policy` file also supplies a default value for the ID of the resource on which the joblet runs. The default value is `.*`, which means all resources are included (see lines 26-29 in [whoami.policy \(page 164\)](#)).

Note that the setting for `resource` is not used in the JDL code but is used to affect which resources on which the joblet run. This occurs because a constraint is specified in `whoami.policy` that restricts the resources that can run this joblet to the current value of the `resource` fact (see line 45 in [whoami.jdl \(page 162\)](#)).

`maxperresource` is another job setting that affects scheduling of the Whoami joblet. The system uses `maxperresource` to determine how many instances of the joblet can run simultaneously on the same resource. In this case, only one instance of the Whoami job can be run on a machine at a time, as specified in lines 39-42 in [whoami.policy \(page 164\)](#).

When facts are referenced in the JDL file, they are prepended with `jobargs.` or `job.` However, when supplied on the command line, this prefix is omitted. JDL files must use an explicit naming convention when it references facts from the different sections of the policy files. For more information on naming conventions for policy facts, see [Section 3.1.1, "PlateSpin Orchestrate Datagrid Filepaths," on page 25](#).

joblet_started_event

When the Whoami joblet is executed on a particular resource it receives a `joblet_started_event`. After this happens, the Whoami joblet decides which command to use to get the current username by checking the value of `resource.os.family` (see lines 50-55 in [whoami.jdl \(page 162\)](#)). This setting is not set in the `whoami.policy`, but instead is available from the PlateSpin Orchestrate system.

After the command to get the current username has been decided, the PlateSpin Orchestrate API class named `Exec` is used to execute the command on the resource where the joblet is running (see lines 62-66 in [whoami.jdl \(page 162\)](#)).

By passing the command to the `Exec setShellCommand` method, the command will be executed by the operating system's default command interpreter. On Microsoft Windows this is `cmd.exe`. On POSIX systems, this is `/bin/sh`. As indicated by lines 64-65 in [whoami.jdl \(page 162\)](#), all standard out and standard errors are written to the job's log file.

To view the log file for the whoami job after it has been run, execute the command `> zos log whoami.`

For more information about using the `zos` command line, see "[The zosadmin Command Line Tool](#)" in the [PlateSpin Orchestrate 2.0 Command Line Reference](#). For more information on running commands using the `Exec` class, see [Exec \(page 219\)](#).

Configure and Run

Execute the following commands to deploy and run `whoami.job`:

- 1 Deploy `notepad.job` into the grid:

```
> zosadmin deploy whoami.job
```

- 2 Display the list of deployed jobs:

```
> zos joblist
```

`whoami` should appear in this list.

- 3 Run the job on one or more resources using the default values for `numJoblets` and `resource`, specified in the `whoami.policy` file:

```
> zos run whoami
```

- 4 Run the job on one or more resources using supplied values for `numJoblets` and `resource`:

```
> zos run whoami numJoblets=10 resource=my_resource_.*
```

Run 10 joblets simultaneously, but only on resources beginning with the name "my_resource_".

NOTE: The value for "resource" is specified using regular expression syntax.

See Also

- ♦ Setting Constraints Using Policies ([Section 4.4, "Policy Management," on page 34](#) and [Chapter 5, "Policy Elements," on page 45](#)).
- ♦ Adding Jobs to Groups During Deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job \(page 168\)](#)).
- ♦ Scheduling multiple instances of a joblet
- ♦ Setting default parameter values using policies
- ♦ Configuring constraints in a policy file
- ♦ Naming conventions for policy facts ([Section 3.1.1, "PlateSpin Orchestrate Datagrid Filepaths," on page 25](#).[Section 3.1.1, "PlateSpin Orchestrate Datagrid Filepaths," on page 25](#))
- ♦ Facts provided by the PlateSpin Orchestrate system that can be referenced within a JDL file
- ♦ Using PlateSpin Orchestrate ("[How Do I Interact with PlateSpin Orchestrate?](#)")
- ♦ Running commands using the `Exec` class.

10.6 Miscellaneous Code-Only Jobs

The following examples demonstrate useful, miscellaneous code-only job concepts:

- ♦ "[factJunction.job](#)" on page 168
- ♦ "[jobargs.job](#)" on page 177

factJunction.job

Demonstrates using fact junctions to retrieve information about objects in the grid relative to another object.

Detail

Each object in the grid has a set of facts which can be read and modified. Some of these facts are special in the sense that their value contains the name of another object that must exist in the grid. These special facts are called fact junctions.

Fact junctions provide a way to reference the facts of one object, using another object as a starting point. For example, all jobs in the grid have a fact named `job.accountinggroup`. The value for `job.accountinggroup` must be the name of a job group currently existing in the grid (the default being the group named `all`). The following JDL code prints the ID of the accounting group for the job named `myJob` without using fact junctions:

```
job = getMatrix().getGridObject(TYPE_JOB, "myJob")
groupName = job.getFact("job.accountinggroup")
group = getMatrix().getGridObject(TYPE_JOBGROUP, groupName)
print "Group ID: " + group.getFact("group.id")
```

Using fact junctions, you can obtain the ID of the accounting group without having to retrieve a reference to the group object first, as follows:

```
job = getMatrix().getGridObject(TYPE_JOB, "myJob")
print "Group ID: " + job.getFact("job.accountinggroup.id")
```

Notice the job `myJob` does not have a fact named `job.accountinggroup.id`. However, it does have a fact named `job.accountinggroup`, which contains the name of an existing job group. This job group has the fact `group.id`, and using fact junctions you can obtain the value of this fact without explicitly reading it off of the job group object itself.

Usage

```
> zosadmin login --user zosadmin Login to server: skate
Please enter current password for 'zosadmin':
Logged into grid on server 'skate'
```

```
> cd /opt/novell/zenworks/zos/server/examples
> zosadmin deploy factJunction.job
factJunction successfully deployed
```

```
> zos login --user zenuser
Please enter current password for 'zenuser':
  Logged into grid as zenuser
```

```
> zos jobinfo --detail factJunction
Jobname/Parameters      Attributes
-----
factJunction           Desc: This is a test job to exercise fact junctions.
```

```
No parameters defined for this job.
```


Description

The files that make up the factJunction job include:

```
factJunction                               # Total: 205 lines
|-- factJunction.jdl                       # 179 lines
`-- factJunction.policy                     # 26 lines
```

factJunction.jdl

```
1 # -----
-----
2 # Copyright 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: factJunction.jdl,v 1.3 2008/02/27 20:49:39 john Exp $
13 # -----
-----
14 #
15 #
16 # This is a test job, but also illustrates all the implemented fact
junctions.
17 # A fact junction is a way to access facts on a 'referenced' object.
18 # E.g. vmhost.resource.*** redirects from the vmhost object through the
19 # junction onto the underlying physical resource object.
20 #
21 # To setup for test, copy job into the 'provisionAdapter' job group.
22 #
23 #
24 #
25 # Add to the 'examples' group on deployment
26 #
27 if __mode__ == "deploy":
28     try:
29         jobgroupname = "examples"
30         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
31         if jobgroup == None:
32             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
33         jobgroup.addMember(__jobname__)
34     except:
35         exc_type, exc_value, exc_traceback = sys.exc_info()
36         print "Error adding %s to %s group: %s %s" % (__jobname__,
jobgroupname, exc_type, exc_value)
37
38
39 class factJunctionJob(Job):
```

```

40
41     def job_started_event(self):
42         m = getMatrix()
43         nptt = "<Not Possible To Test>"
44
45         # Setup test environment
46         user = m.getGridObject(TYPE_USER, "test")
47         if (user == None):
48             user = m.createGridObject(TYPE_USER, "test")
49         user.setArrayFact("user.privilegedjobgroups", ["all"])
50
51         repository = m.getGridObject(TYPE_REPOSITORY, "test")
52         if (repository == None):
53             repository = m.createGridObject(TYPE_REPOSITORY, "test")
54         repository.setArrayFact("repository.provisioner.jobs",
["factJunction"])
55
56         node = m.getGridObject(TYPE_RESOURCE, "test")
57         if (node == None):
58             node = m.createGridObject(TYPE_RESOURCE, "test")
59
60         vm = m.getGridObject(TYPE_RESOURCE, "vmtest")
61         if (vm == None):
62             vm = m.createResource("vmtest", ResourceInfo.TYPE_VM_INSTANCE)
63         vm.setFact("resource.provisioner.job", "factJunction")
64         vm.setFact("resource.vm.repository", "test")
65         vm.setFact("resource.provisioner.recommendedhost", "test_test")
66
67         vmt = m.getGridObject(TYPE_RESOURCE, "vmttest")
68         if (vmt == None):
69             vmt = m.createResource("vmttest",
ResourceInfo.TYPE_VM_TEMPLATE)
70         vmt.setFact("resource.provisioner.job", "factJunction")
71         vmt.setFact("resource.vm.repository", "test")
72
73         try:
74             vmhost = node.getVmHost("test")
75         except:
76             vmhost = node.createVmHost("test")
77         vmhost.setFact("vmhost.provisioner.job", "factJunction")
78         vmhost.setArrayFact("vmhost.repositories", ["test"])
79         vmhost.setArrayFact("vmhost.vm.available.groups", ["all"])
80
81         job = m.getGridObject(TYPE_JOB, "factJunction")
82
83         # Test junctions
84
85         print
86         print "Testing User fact junctions (3):"
87         r = user.getFact("user.accountinggroup.id")
88         print "1. user.accountinggroup.id = %s" % r
89         # Array junctions
90         r = user.getFact("user.privilegedjobgroups[all].id")
91         print "2. user.privilegedjobgroups[all].id = %s" % r
92         r = user.getFact("user.groups[all].jobcount")
93         print "3. user.groups[all].jobcount = %s" % r
94
95         print
96

```

```

97     print "Testing Job fact junctions (3):"
98     r = job.getFact("job.accountinggroup.id")
99     print "1. job.accountinggroup.id = %s" % r
100    r = job.getFact("job.resourcegroup.id")
101    print "2. job.resourcegroup.id = %s" % r
102    # Array junctions
103    r = job.getFact("job.groups[all].jobinstances.total")
104    print "3. job.groups[all].jobinstances.total = %s" % r
105
106
107    print
108    print "Testing VmHost fact junctions (7):"
109    r = vmhost.getFact("vmhost.resource.id")
110    print "1. vmhost.resource.id = %s" % r
111    r = vmhost.getFact("vmhost.accountinggroup.id")
112    print "2. vmhost.accountinggroup.id = %s" % r
113    r = vmhost.getFact("vmhost.provisioner.job.id")
114    print "3. vmhost.provisioner.job.id = %s" % r
115    # Array junctions
116    r = vmhost.getFact("vmhost.groups[all].vmcount")
117    print "4. vmhost.groups[all].vmcount = %s" % r
118    r = vmhost.getFact("vmhost.repositories[test].id")
119    print "5. vmhost.repositories[test].id = %s" % r
120    r = vmhost.getFact("vmhost.vm.available.groups[all].id")
121    print "6. vmhost.vm.available.groups[all].id = %s" % r
122    #r = vmhost.getFact("vmhost.vm.instanceids[vmtest].id")
123    r = npptt
124    print "7. vmhost.vm.instanceids.[vmtest].id = %s" % r
125
126
127    print
128    print "Testing Resource fact junctions (9):"
129    r = vm.getFact("resource.provisioner.job.id")
130    print "1. resource.provisioner.job.id = %s" % r
131    r = vm.getFact("resource.vm.repository.id")
132    print "2. resource.vm.repository.id = %s" % r
133    r = vm.getFact("resource.provisioner.recommendedhost.id")
134    print "3. resource.provisioner.recommendedhost.id = %s" % r
135    #r = vm.getFact("resource.provision.vmhost.id")
136    r = npptt
137    print "4. resource.provision.vmhost.id = %s" % r
138    #r = vm.getFact("resource.provision.template.id")
139    r = npptt
140    print "5. resource.provision.template.id = %s" % r
141    # Array junctions
142    r = vm.getFact("resource.groups[all].loadaverage")
143    print "6. resource.groups[all].loadaverage = %s" % r
144    r = node.getFact("resource.vmhosts[test_test].id")
145    print "7. resource.vmhosts[test_test].id = %s" % r
146    r = node.getFact("resource.repositories[test].id")
147    print "8. resource.repositories[test].id = %s" % r
148    #r = vmt.getFact("resource.provisioner.instances[vmtestest_2].id")
149    r = npptt
150    print "9. resource.provisioner.instances[vmtestest_2].id = %s" % r
151
152
153    print
154    print "Testing Repository fact junctions (4):"
155    r = repository.getFact("repository.groups[all].id")

```

```

156     print "1. repository.groups[all].id = %s" % r
157     r = repository.getFact("repository.vmimages[vmtest].id")
158     print "2. repository.vmimages[vmtest].id = %s" % r
159     r = repository.getFact("repository.vmhosts[test_test].id")
160     print "3. repository.vmhosts[test_test].id = %s" % r
161     r =
repository.getFact("repository.provisioner.jobs[factJunction].id")
162     print "4. repository.provisioner.jobs[factJunction].id = %s" % r
163
164
165     print
166     print "Testing multiple junctions (1):"
167     r =
repository.getFact("repository.vmhosts[test_test].resource.repositories[test]
.vmhosts[test_test].groups[all].id")
168     print "1.
repository.vmhosts[test_test].resource.repositories[test].vmhosts[test_test].
groups[all].id = %s" % r
169
170     # Now make sure they are all accessible by the joblet...
171     #self.schedule(factJunctionJoblet, {})
172
173
174 class factJunctionJoblet(Joblet):
175
176     def joblet_started_event(self):
177         # TODO
178         time.sleep(sleeptime)
179

```

factJunction.policy

The description fact displays the commands x, y, z ...

```

1  <!--
2
=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT.  NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
=====
=
13 * $Id: factJunction.policy,v 1.2 2008/02/27 20:49:39 john Exp $
14
=====
=

```

```

15  -->
16
17  <policy>
18
19    <job>
20      <fact name="description"
21        type="String"
22        value="This is a test job to exercise fact junctions." />
23    </job>
24
25  </policy>
26

```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

UserInfo

A representation of a User grid object. This class provides accessors and setters for User facts.

RepositoryInfo

A representation of a Repository grid object. This class provides accessors and setters for Repository facts. To script the creation of Repository objects, see [MatrixInfo \(page 235\)](#).

ResourceInfo

A representation of a Resource Grid Object. This class inherits the base fact operations from `GridObjectInfo` and adds the provisioning operations for provisionable resources such as virtual machines. See [MatrixInfo \(page 235\)](#) for how to script creation of Resource objects.

JobInfo

A representation of a deployed Job. The factset available on the JobInfo class is the aggregation of the Job's policy and policies on the groups the Job is a member of. This includes the `job.*` and `jobargs.*` fact namespaces.

Job Details

The FactJunction job performs its work by handling the following events:

- ♦ [“zosadmin deploy” on page 174](#)
- ♦ [“job_started_event” on page 174](#)
- ♦ [“joblet_started_event” on page 174](#)

zosadmin deploy

Deploying FactJunction job is performed by lines 27-36 of factJunction.jdl. When jobs are deployed into the grid, they can optionally be placed in groups for organization and easy reference. In this case, the FactJunction job will be added to the group named “examples”, and will show up in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples
```

For a general overview of how jobs are added to groups during deployment, see [“Walkthrough: Deploying a Sample Job”](#) in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

job_started_event

When the FactJunction job receives a `job_started_event`, it gets a reference to the `MatrixInfo` object (line 42), which allows it to obtain references to other objects in the grid, such as Users, Resources, Jobs, etc. (see lines 46, 51, 56, 60, 67 and 81 in [“factJunction.jdl” on page 169](#)). If these objects don't exist in the grid, they are immediately created so they can be used later on (see lines 48, 53, 58, 62, 69 and 76).

After references exist for the various objects in the grid, values for other objects are printed out using the fact junctions that exist on each object (see lines 83-168 in [factJunction.jdl \(page 169\)](#)).

There are several instances where the factJunction job uses “array notation” to handle fact junctions that contain multiple values (see lines 90, 92, 103, 116, 120, 122, 124, 142, 144, 146, 156, 158, 160 and 167 in [factJunction.jdl \(page 169\)](#)). As previously explained, fact junctions are special facts because their value contains the name of another object that must exist in the grid. However, fact junctions don't always contain a single name. Some fact junctions allow for an array of names to be specified. For example, the value for the fact “job.groups” is supplied as a String array.

In this case, the fact junction can be refined using array notation, which allows for the selection of one of the values. For example, the following code retrieves the ID of the group named “myGroup”, which is one of the groups the given job is a member of:

```
job.getFact("job.groups[myGroup].id")
```

joblet_started_event

The FactJunction job only illustrates using fact junctions to retrieve information about objects in the grid. Therefore, no work is performed on the resource by the FactJunction joblet.

Configure and Run

To run this example, you must have PlateSpin Orchestrate installed and configured properly. No agents on separate resources are required. You also must be logged into your Orchestrate server before you run `zosadmin` or `zos` commands.

Execute the following commands to deploy and run `factJunction.job`:

1 Deploy `factJunction.job` into the grid:

```
> zosadmin deploy factJunction.job
```

2 Display the list of deployed jobs:

```
> zos joblist
```

`factJunction` should appear in this list.

3 Run the `FactJunction` job, and view the results:

```
> zos run factJunction
JobID: zenuser.factJunction.421
> zos log factJunction
> zos status zenuser.factJunction.421
Completed
```

```
> zos log factJunction
```

4 Testing User fact junctions:

1. `user.accountinggroup.id = all`
2. `user.privilegedjobgroups[all].id = all`
3. `user.groups[all].jobcount = 147`

5 Testing Job fact junctions:

1. `job.accountinggroup.id = all`
2. `job.resourcegroup.id = all`
3. `job.groups[all].jobinstances.total = 1`

6 Testing VmHost fact junctions:

1. `vmhost.resource.id = test`
2. `vmhost.accountinggroup.id = all`
3. `vmhost.provisioner.job.id = factJunction`
4. `vmhost.groups[all].vmcount = 0`
5. `vmhost.repositories[test].id = test`
6. `vmhost.vm.available.groups[all].id = all`
7. `vmhost.vm.instanceids.[vmtest].id = <Not Possible To Test>`

7 Testing Resource fact junctions:

1. `resource.provisioner.job.id = factJunction`
2. `resource.vm.repository.id = test`
3. `resource.provisioner.recommendedhost.id = test_test`
4. `resource.provision.vmhost.id = <Not Possible To Test>`
5. `resource.provision.template.id = <Not Possible To Test>`

6. `resource.groups[all].loadaverage = 0.03666666666666667`
7. `resource.vmhosts[test_test].id = test_test`
8. `resource.repositories[test].id = test`
9. `resource.provisioner.instances[vmtest_2].id = <Not Possible To Test>`

8 Testing Repository fact junctions:

1. `repository.groups[all].id = all`
2. `repository.vmimages[vmtest].id = vmtest`
3. `repository.vmhosts[test_test].id = test_test`
4. `repository.provisioner.jobs[factJunction].id = factJunction`

9 Testing multiple junctions

1. `repository.vmhosts[test_test].resource.repositories[test].vmhosts[test_test].groups[all].id = all`

See Also

- ♦ Adding jobs to groups during deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job \(page 168\)](#)).
- ♦ View the list of fact junctions available for each object type in a PlateSpin Orchestrate grid
- ♦ Using array notation to refine multi-valued fact junctions
- ♦ Using PlateSpin Orchestrate (“[How Do I Interact with PlateSpin Orchestrate?](#)”)

jobargs.job

Demonstrates the usage of the various argument types that jobs can accept. These types are integer, Real, Boolean, String, Time, Date, List, Dictionary, and Array (which can contain the types Integer, Real, Boolean, Time, Date, String). For more information about how to define job arguments, and specify their values on the command line, see [Section 7.7, “Working with Facts and Constraints,” on page 55](#).

Usage

```
> zosadmin login --user zosadmin Login to server: skate
Please enter current password for 'zosadmin':
Logged into grid on server 'skate'
```

```
> cd /opt/novell/zenworks/zos/server/examples
> zosadmin deploy jobargs.job
jobargs successfully deployed
```

```
> zos login --user zenuser Please enter current password for 'zenuser':
Logged into grid as zenuser
```

```
> zos jobinfo --detail jobargs
```

```
Jobname/Parameters  Attributes
-----
```

```
jobargs             Desc: This example job tests all fact types.
```

```
  TimeArgReq        Desc: Required Time arg test
                    Type: Time
                    Default: None! Value must be specified
```

```
  IntegerArg        Desc: Integer arg test
                    Type: Integer
                    Default: 1
```

```
  DateArg           Desc: Date arg test
                    Type: Date
                    Default: Wed Apr 05 09:00:00 EDT 2006
```

```
  RealArgReq        Desc: Required Real arg test
                    Type: Real
                    Default: None! Value must be specified
```

```
  IntegerArrayArg   Desc: Integer[] arg test
                    Type: Integer[]
                    Default: [100,200,300]
```

```
  RealArrayArg      Desc: Real[] arg test
                    Type: Real[]
                    Default: [1.23,3.456,7.0]
```

```
  DictArg           Desc: Dictionary arg test
                    Type: Dictionary
                    Default: {name=moe, dob=Fri Jan 02 00:00:00 EST 1970,
                               age=35}
```

```
  DateArrayArg      Desc: Date[] arg test
```

```

                                Type: Date[]
                                Default: [Wed Jul 08 22:00:00 EDT 2009,Thu Jan 02 00:01:00
                                           EST 2003]

StringArgReq      Desc: Required String arg test
                  Type: String
                  Default: None! Value must be specified

TimeArrayArg     Desc: Time[] arg test
                  Type: Time[]
                  Default: [79200000,60000]

StringArrayArg   Desc: String[] arg test
                  Type: String[]
                  Default: [abc,def,ghi jkl]

TimeArg          Desc: Time arg test
                  Type: Time
                  Default: 32400000

BooleanArgReq    Desc: Required Boolean arg test
                  Type: Boolean
                  Default: None! Value must be specified

BooleanArg       Desc: Boolean arg test
                  Type: Boolean
                  Default: true

IntegerArgReq    Desc: Required Integer arg test
                  Type: Integer
                  Default: None! Value must be specified

StringArg        Desc: String arg test
                  Type: String
                  Default: Hello World

BooleanArrayArg  Desc: Boolean[] arg test
                  Type: Boolean[]
                  Default: [true,false,true]

RealArg          Desc: Real arg test
                  Type: Real
                  Default: 3.1415

ListArg          Desc: List arg test
                  Type: List
                  Default: [abc, d, efghij]

DateArgReq       Desc: Required Date arg test
                  Type: Date
                  Default: None! Value must be specified

```

Description

The files that make up the Jobargs job include:

```

jobargs.job # Total: 254 lines
|-- jobargs.jdl # 77 lines
`-- jobargs.policy # 177 lines

```

jobargs.jdl

```

1 # -----
-----
2 # Copyright © 2008 Novell, Inc. All Rights Reserved.
3 #
4 # NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
5 # WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY ,
6 # FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. NOVELL, THE
AUTHORS
7 # OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
8 # FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
9 # TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
10 # OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
11 # -----
-----
12 # $Id: jobargs.jdl,v 1.3 2008/02/27 20:50:31 john Exp $
13 # -----
-----
14
15 """
16 Example job showing all available job argument types.
17
18 Example cmd line to run job:
19     zos run jobargs TimeArgReq="12:01:02" RealArgReq="3.14" IntegerArgR
eq="123" StringArgReq="foo" BooleanArgReq="true" ListArg="hi,mom"
20 """
21
22 import time
23
24 #
25 # Add to the 'examples' group on deployment
26 #
27 if __mode__ == "deploy":
28     try:
29         jobgroupname = "examples"
30         jobgroup = getMatrix().getGroup(TYPE_JOB, jobgroupname)
31         if jobgroup == None:
32             jobgroup = getMatrix().createGroup(TYPE_JOB, jobgroupname)
33         jobgroup.addMember(__jobname__)
34     except:
35         exc_type, exc_value, exc_traceback = sys.exc_info()
36         print "Error adding %s to %s group: %s %s" % (__jobname__, jobgr
oupname, exc_type, exc_value)
37
38
39 class jobargs(Job):
40
41     def job_started_event(self):
42
43         jobid = self.getFact("jobinstance.id")

```

```

44         print "*****Dumping %s JobInstance jobargs facts*****" % (jobi d)
45         keys = self.getFactNames()
46         keys.sort()
47         for s in keys:
48             if s.startswith("jobargs"):
49                 v = self.getFact(s)
50                 ty = type(v)
51
52                 if str(ty).endswith("Dictionary"):
53                     self.dump_dict(s,v)
54                 else:
55                     if s.endswith("TimeArg") or s.endswith("TimeArgReq "):
56                         v = time.ctime(v/1000)
57
58                         print "%s %s %s" % (s,type(v),str(v))
59         print "*****End %s dump*****" % (jobid)
60
61         #self.schedule(jobargsJoblet)
62
63     def dump_dict(self,name,dict):
64         print "Dict: %s" % (name)
65         keys = dict.keys()
66         for k in keys:
67             v = dict[k]
68             ty = type(v)
69             if k == "dob":
70                 v = time.ctime(v/1000)
71             print "    %s %s %s" % (k,ty,str(v))
72
73
74     class jobargsJoblet(Joblet):

```

jobargs.policy

```

1  <!--
2
=====
=
3  * Copyright © 2008 Novell, Inc. All Rights Reserved.
4  *
5  * NOVELL PROVIDES THE SOFTWARE "AS IS," WITHOUT ANY EXPRESS OR IMPLIED
6  * WARRANTY, INCLUDING WITHOUT THE IMPLIED WARRANTIES OF MERCHANTABILITY,
7  * FITNESS FOR A PARTICULAR PURPOSE, AND NON INFRINGEMENT.  NOVELL, THE
AUTHORS
8  * OF THE SOFTWARE, AND THE OWNERS OF COPYRIGHT IN THE SOFTWARE ARE NOT
LIABLE
9  * FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF
CONTRACT,
10 * TORT, OR OTHERWISE, ARISING FROM, OUT OF, OR IN CONNECTION WITH THE
SOFTWARE
11 * OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
12
=====
=
13 * $Id: jobargs.policy,v 1.2 2008/02/27 20:50:31 john Exp $
14
=====
=

```

```

15  -->
16
17 <policy>
18
19   <jobargs>
20
21     <!-- Optional job args -->
22     <fact name="OptionalDateArg"
23         description="Optional Date Arg"
24         type="Date"
25         value="1/2/06 12:01 PM"/>
26
27     <fact name="OptionalTimeArg"
28         description="Optional Time Arg"
29         type="Time"
30         value="12:01 PM"/>
31
32     <fact name="OptionalRealArg"
33         description="Optional Real Arg"
34         type="Real"
35         value="3.14" />
36
37     <fact name="OptionalIntegerArg"
38         description="Optional Integer Arg"
39         type="Integer"
40         value="123" />
41
42     <fact name="OptionalStringArg"
43         description="Optional String Arg"
44         type="String"
45         value="foo" />
46
47     <fact name="OptionalString2ArgAsTag"
48         description="Optional String Arg as tag">
49       <string>bar</string>
50     </fact>
51
52     <fact name="OptionalString3ArgAsCDATA"
53         description="Optional String Arg as CDATA">
54       <string>
55         <![CDATA[this text is part of
56 a multi-line cdata section containing
57 xml <html>test</html>
58 <eq fact="foo.bar" value="qwerty" />
59 cool!
60 ]]>
61       </string>
62     </fact>
63
64     <fact name="OptionalBooleanArg"
65         description="Optional Boolean Arg"
66         type="Boolean"
67         value="true" />
68
69     <fact name="OptionalListArg">
70       <list>
71         <listelement value="hi" type="String" />
72         <listelement value="mom" />
73         <listelement value="42" type="Integer" />

```

```

74     </list>
75 </fact>
76
77 <fact name="OptionalDictArg">
78   <dictionary>
79     <dictelement key="name" type="String" value="joe" />
80     <dictelement key="date" type="Date" value="4/15/06" />
81     <dictelement key="time" type="Time" value="3:30 AM" />
82     <dictelement key="age" type="Integer" value="12" />
83   </dictionary>
84 </fact>
85
86 <fact name="OptionalDateArray">
87   <array>
88     <date>1/2/06 12:01 PM</date>
89     <date>1/3/06 12:02 PM</date>
90     <date>1/4/06</date>
91   </array>
92 </fact>
93 <fact name="OptionalTimeArray">
94   <array>
95     <time>12:01 PM</time>
96     <time>12:02 PM</time>
97   </array>
98 </fact>
99 <fact name="OptionalRealArray">
100   <array>
101     <real>1.1</real>
102     <real>2.2</real>
103   </array>
104 </fact>
105 <fact name="OptionalIntegerArray">
106   <array>
107     <integer>1</integer>
108     <integer>2</integer>
109   </array>
110 </fact>
111 <fact name="OptionalStringArray">
112   <array>
113     <string>string1</string>
114     <string>string2</string>
115   </array>
116 </fact>
117 <!-- Arrays of dictionary or list not currently supported
118   <fact name="OptionalDictionaryArray">
119     <array>
120       <dictionary>
121         <dictelement key="name" type="String" value="joe" />
122       </dictionary>
123     </array>
124   </fact>
125 -->
126
127 <!-- Required job args -->
128 <fact name="RequiredDateArg" type="Date" />
129 <fact name="RequiredTimeArg" type="Time" />
130 <fact name="RequiredRealArg" type="Real" />
131 <fact name="RequiredIntegerArg" type="Integer" />
132 <fact name="RequiredStringArg" type="String" />

```

```

133     <fact name="RequiredBooleanArg" type="Boolean" />
134 <!-- XXX Ooops, not currently supported without value!
135     <fact name="RequiredListArg" type="list" />
136     <fact name="RequiredDictArg" type="dictionary" />
137     <fact name="RequiredStringArray" type="string">
138         <array />
139     </fact>
140 -->
141
142     <!-- Invisible job args -->
143     <fact name="InvisibleDateArg" type="Date" value="1/2/06 12:01 PM"
visible="False" />
144     <fact name="InvisibleTimeArg" type="Time" value="12:01 PM"
visible="False" />
145     <fact name="InvisibleRealArg" type="Real" value="3.14"
visible="False" />
146     <fact name="InvisibleIntegerArg" type="Integer" value="123"
visible="False" />
147     <fact name="InvisibleStringArg" type="String" value="foo"
visible="False" />
148     <fact name="InvisibleString2Arg" visible="False" >
149         <string>bar</string>
150     </fact>
151     <fact name="InvisibleBooleanArg" type="Boolean" value="true"
visible="False" />
152     <fact name="InvisibleListArg" visible="False">
153         <list>
154             <listelement value="hi" type="String" />
155             <listelement value="mom" />
156             <listelement value="42" type="integer" />
157         </list>
158     </fact>
159     <fact name="InvisibleDictArg" visible="False">
160         <dictionary>
161             <dictelement key="name" type="String" value="joe" />
162             <dictelement key="date" type="Date" value="4/15/06" />
163             <dictelement key="time" type="Time" value="3:30 AM" />
164             <dictelement key="age" type="Integer" value="12" />
165         </dictionary>
166     </fact>
167
168 </jobargs>
169
170 <job>
171     <fact name="description"
172         type="String"
173         value="This example job tests all fact types." />
174 </job>
175
176 </policy>
177

```

Schedule File (optional)

jobargs.sched

```
1 <schedule name="jobargs" description="Run jobargs" active="true">
2 <runjob job="jobargs" user="labuser" priority="medium" />
3 <triggers>
4 <trigger name="trigger1" />
5 <trigger name="trigger2" />
6 </triggers>
7 </schedule>
```

Classes and Methods

Definitions:

Job

A representation of a running job instance.

Joblet

Defines execution on the resource.

MatrixInfo

A representation of the matrix grid object, which provides operations for retrieving and creating grid objects in the system. MatrixInfo is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

GroupInfo

A representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group.

Job Details

The Jobargs job performs its work by handling the following events:

- ♦ “zosadmin deploy” on page 184
- ♦ “job_started_event” on page 185
- ♦ “joblet_started_event” on page 185

zosadmin deploy

In `jobargs.jdl` (page 179), lines 27-36 deploy the job into the grid. After jobs are deployed into the grid, they can optionally be placed in groups for organization and easy reference. In this case, the jobargs job will be added to the group named “examples”, and will show up in the PlateSpin Orchestrate Development Client in the Explorer view at the location:

```
/ZOS/YOUR_GRID/Jobs/examples
```

For a general overview of how jobs are added to groups during deployment, see “[Walkthrough: Deploying a Sample Job](#)” in the *PlateSpin Orchestrate 2.0 Installation and Configuration Guide*.

job_started_event

After the Jobargs job receives a `job_started_event`, it gets a list of all the facts available to it, as shown in line 45 of `jobargs.jdl` (page 179). This list is sorted, filtered according to whether or not it's a jobarg fact, and then enumerated (lines 46-58). Each jobarg fact is printed in a "name type value" format. When the complex Dictionary type is encountered (line 52), a separate method is used to print the values for all the key-value pairs (lines 63-71).

The list of optional and required arguments for this Jobargs example are available as facts within the `<jobargs>` section (see lines 19-168 in `jobargs.policy` (page 180)).

For more information about defining job arguments and their types, see [Chapter 5, "Policy Elements,"](#) on page 45 and [Section 4.4, "Policy Management,"](#) on page 34.

joblet_started_event

The Jobargs job only illustrates passing job arguments to a job. Therefore, no work is performed on the resource by the `jobargsJoblet`.

Configure and Run

To run this example, you must have PlateSpin Orchestrate installed and configured properly. No agents on separate resources are required. You also must be logged into your Orchestrate Server before you run `zosadmin` or `zos` commands.

Execute the following commands to deploy and run `jobargs.job`:

- 1 Deploy `jobargs.job` into the grid:

```
> zosadmin deploy jobarg.job
```

NOTE: Run `zosadmin login` to log in for `zos` administration.

- 2 Display the list of deployed jobs:

```
> zos joblist
```

jobargs should appear in this list.

NOTE: Run `zos login` to run `zos` client jobs.

- 3 Display the list of optional and required arguments for this job:

```
> zos jobinfo jobargs
```

- 4 Run the `jobargs` job and view the results.

NOTE: You must supply values for `TimeArgReq`, `RealArgReq`, `StringArgReq`, `BooleanArgReq`, `IntegerArgReq`, and `DateArgReq` as follows (see [jobargs.policy](#) (page 180) for the full list of arguments that can be specified):

```
> zos run jobargs TimeArgReq=12:01:02 RealArgReq=3.14 StringArgReq=Hello  
BooleanArgReq=True IntegerArgReq=42 DateArgReq="04/05/07 7:45 AM"
```

```
> zos log jobargs
```

See Also

- ♦ Adding Jobs to Groups During Deployment (see how the JDL code can print the ID of group of jobs in [factJunction.job \(page 168\)](#)).
- ♦ Defining job arguments and their types
- ♦ Using PlateSpin Orchestrate (“[How Do I Interact with PlateSpin Orchestrate?](#)”)

PlateSpin Orchestrate Client SDK

A

This section provides the reference information for the Java* classes used by the PlateSpin® Orchestrate Client SDK

- ◆ [Section A.1, “Constraint Package,” on page 187](#)
- ◆ [Section A.2, “Datagrid Package,” on page 191](#)
- ◆ [Section A.3, “Grid Package,” on page 192](#)
- ◆ [Section A.4, “TLS Package,” on page 199](#)
- ◆ [Section A.5, “Toolkit Package,” on page 199](#)

A.1 Constraint Package

The following Java files form the interfaces and exceptions for the PlateSpin Orchestrate constraint grid structure:

- ◆ [Section A.1.1, “AndConstraint,” on page 187](#)
- ◆ [Section A.1.2, “BetweenConstraint,” on page 188](#)
- ◆ [Section A.1.3, “BinaryConstraint,” on page 188](#)
- ◆ [Section A.1.4, “Constraint,” on page 188](#)
- ◆ [Section A.1.5, “ContainerConstraint,” on page 188](#)
- ◆ [Section A.1.6, “ContainsConstraint,” on page 188](#)
- ◆ [Section A.1.7, “DefinedConstraint,” on page 188](#)
- ◆ [Section A.1.8, “EqConstraint,” on page 189](#)
- ◆ [Section A.1.9, “GeConstraint,” on page 189](#)
- ◆ [Section A.1.10, “GtConstraint,” on page 189](#)
- ◆ [Section A.1.11, “IfConstraint,” on page 189](#)
- ◆ [Section A.1.12, “LeConstraint,” on page 189](#)
- ◆ [Section A.1.13, “LtConstraint,” on page 189](#)
- ◆ [Section A.1.14, “NeConstraint,” on page 189](#)
- ◆ [Section A.1.15, “NotConstraint,” on page 190](#)
- ◆ [Section A.1.16, “OperatorConstraint,” on page 190](#)
- ◆ [Section A.1.17, “OrConstraint,” on page 190](#)
- ◆ [Section A.1.18, “TypedConstraint,” on page 190](#)
- ◆ [Section A.1.19, “UndefinedConstraint,” on page 190](#)
- ◆ [Section A.1.20, “ConstraintException,” on page 190](#)

A.1.1 AndConstraint

Perform a logical and-ing of all child constraints. This is a no-op if this constraint contains no children.

For complete documentation of the class, see [AndConstraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/AndConstraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/AndConstraint.html).

A.1.2 BetweenConstraint

Binary Operator Constraints that have both a left and right side.

For complete documentation of the class, see [BetweenConstraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/BetweenConstraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/BetweenConstraint.html).

A.1.3 BinaryConstraint

Binary Operator Constraints that have both a left and right side.

For complete documentation of the class, see [BinaryConstraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/BinaryConstraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/BinaryConstraint.html).

A.1.4 Constraint

Basic Constraint interface that allows traversal and evaluation of a constraint tree.

For complete documentation of the class, see [Constraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/Constraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/Constraint.html).

A.1.5 ContainerConstraint

Container constraints that perform logical aggregation operations on contained constraints.

For complete documentation of the class, see [ContainerConstraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/ContainerConstraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/ContainerConstraint.html).

A.1.6 ContainsConstraint

Performs a simple set operation that returns true if the right side of the operation is found in the value set of the left side.

For complete documentation of the class, see [ContainsConstraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/ContainsConstraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/ContainsConstraint.html).

A.1.7 DefinedConstraint

Evaluates to true only if the left side fact is defined in the match context.

For complete documentation of the class, see [DefinedConstraint \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/DefinedConstraint.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/DefinedConstraint.html).

A.1.8 EqConstraint

Performs a equality constraint operation.

For complete documentation of the class, see [EqConstraint \(http://www.novell.com/documentation/zen_orchestrator13/reference/javadoc/com/novell/zos/constraint/EqConstraint.html\)](http://www.novell.com/documentation/zen_orchestrator13/reference/javadoc/com/novell/zos/constraint/EqConstraint.html).

A.1.9 GeConstraint

Performs a 'greater than or equal to' constraint operation.

For complete documentation of the class, see [GeConstraint \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/GeConstraint.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/GeConstraint.html).

A.1.10 GtConstraint

Performs a 'greater than' constraint operation.

For complete documentation of the class, see [GtConstraint \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/GtConstraint.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/GtConstraint.html).

A.1.11 IfConstraint

Perform a conditional if,then,else block.

For complete documentation of the class, see [IfConstraint \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/IfConstraint.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/IfConstraint.html).

A.1.12 LeConstraint

Performs a 'less than or equal to' constraint operation.

For complete documentation of the class, see [LeConstraint \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/LeConstraint.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/LeConstraint.html).

A.1.13 LtConstraint

Performs a 'less than' constraint operation.

For complete documentation of the class, see [LtConstraint \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/LtConstraint.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/LtConstraint.html).

A.1.14 NeConstraint

Performs a not equal constraint operation.

For complete documentation of the class, see [NeConstraint \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/NeConstraint.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/constraint/NeConstraint.html).

A.1.15 NotConstraint

Perform a logical not operation of all the child constraints.

For complete documentation of the class, see [NotConstraint](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/NotConstraint.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/NotConstraint.html).

A.1.16 OperatorConstraint

Operator constraints that perform comparison operation on facts.

For complete documentation of the class, see [OperatorConstraint](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/OperatorConstraint.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/OperatorConstraint.html).

A.1.17 OrConstraint

Perform a logical or-ing operation of all the child constraints.

For complete documentation of the class, see [OrConstraint](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/OrConstraint.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/OrConstraint.html).

A.1.18 TypedConstraint

Typed constraint must only be used as the outermost wrapper when it is necessary to override the default constraint type of 'resource'.

For complete documentation of the class, see [TypedConstraint](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/TypedConstraint.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/TypedConstraint.html).

A.1.19 UndefinedConstraint

Evaluates to true only if the left side fact is not defined in the match context.

For complete documentation of the class, see [UndefinedConstraint](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/UndefinedConstraint.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/UndefinedConstraint.html).

A.1.20 ConstraintException

For exceptions that occur in parsing or executing constraints.

For complete documentation of the class, see [Constraint exception](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/ConstraintException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/constraint/ConstraintException.html).

A.2 Datagrid Package

The following Java files form the classes, interfaces, and exceptions for the PlateSpin Orchestrate datagrid structure:

- ♦ [Section A.2.1, “GridFile,” on page 191](#)
- ♦ [Section A.2.2, “GridFileFilter,” on page 191](#)
- ♦ [Section A.2.3, “GridFileNameFilter,” on page 191](#)
- ♦ [Section A.2.4, “DGLogger,” on page 191](#)
- ♦ [Section A.2.5, “DataGridException,” on page 191](#)
- ♦ [Section A.2.6, “DataGridNotAvailableException,” on page 192](#)
- ♦ [Section A.2.7, “GridFile.CancelException,” on page 192](#)

A.2.1 GridFile

Specifies the PlateSpin Orchestrate datagrid interface for individual files and directories.

For complete documentation of the class, see [GridFile](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFile.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFile.html).

A.2.2 GridFileFilter

Filter for accepting/rejecting file names in a directory list.

For complete documentation of the class, see [GridFileFilter](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFileFilter.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFileFilter.html).

A.2.3 GridFileNameFilter

Filter for accepting/rejecting file names in a directory list.

For complete documentation of the class, see [GridFileNameFilter](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFileNameFilter.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFileNameFilter.html).

A.2.4 DGLogger

Definitions of the DataGrid Logger options used for multicast.

For complete documentation of the class, see [DGLogger](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/DGLogger.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/DGLogger.html).

A.2.5 DataGridException

General exception class for datagrid errors.

For complete documentation of the class, see [DataGridException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/DataGridException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/DataGridException.html).

A.2.6 DataGridNotAvailableException

Exception thrown if the datagrid cannot be reached due to a network error.

For complete documentation of the class, see [DataGridNotAvailableException \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/DataGridNotAvailableException.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/DataGridNotAvailableException.html).

A.2.7 GridFile.CancelException

Exception thrown by cancelled requests.

For complete documentation of the class, see [GridFile.CancelException \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFile.CancelException.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/dataGrid/GridFile.CancelException.html).

A.3 Grid Package

The Java classes included in the Grid package form the basis of the PlateSpin Orchestrate infrastructure. For complete documentation of each class, click on the links to access the online documentation Javadoc.

- ◆ [Section A.3.1, “AgentListener,” on page 193](#)
- ◆ [Section A.3.2, “ClientAgent,” on page 193](#)
- ◆ [Section A.3.3, “Credential,” on page 193](#)
- ◆ [Section A.3.4, “Fact,” on page 193](#)
- ◆ [Section A.3.5, “FactSet,” on page 194](#)
- ◆ [Section A.3.6, “GridObjectInfo,” on page 194](#)
- ◆ [Section A.3.7, “ID,” on page 194](#)
- ◆ [Section A.3.8, “JobInfo,” on page 194](#)
- ◆ [Section A.3.9, “Message,” on page 194](#)
- ◆ [Section A.3.10, “Message.Ack,” on page 194](#)
- ◆ [Section A.3.11, “Message.AuthFailure,” on page 194](#)
- ◆ [Section A.3.12, “Message.ClientResponseMessage,” on page 195](#)
- ◆ [Section A.3.13, “Message.ConnectionID,” on page 195](#)
- ◆ [Section A.3.14, “Message.Event,” on page 195](#)
- ◆ [Section A.3.15, “Message.GetGridObjects,” on page 195](#)
- ◆ [Section A.3.16, “Message.GridObjects,” on page 195](#)
- ◆ [Section A.3.17, “Message.JobAccepted,” on page 195](#)
- ◆ [Section A.3.18, “Message.JobError,” on page 195](#)
- ◆ [Section A.3.19, “Message.JobFinished,” on page 196](#)
- ◆ [Section A.3.20, “Message.JobIdEvent,” on page 196](#)
- ◆ [Section A.3.21, “Message.JobInfo,” on page 196](#)
- ◆ [Section A.3.22, “Message.Jobs,” on page 196](#)

- ◆ Section A.3.23, “Message.JobStarted,” on page 196
- ◆ Section A.3.24, “Message.JobStatus,” on page 196
- ◆ Section A.3.25, “Message.LoginFailed,” on page 196
- ◆ Section A.3.26, “Message.LoginSuccess,” on page 197
- ◆ Section A.3.27, “Message.LogoutAck,” on page 197
- ◆ Section A.3.28, “Message.RunningJobs,” on page 197
- ◆ Section A.3.29, “Message.ServerStatus,” on page 197
- ◆ Section A.3.30, “Node,” on page 197
- ◆ Section A.3.31, “Priority,” on page 197
- ◆ Section A.3.32, “WorkflowInfo,” on page 197
- ◆ Section A.3.33, “ClientOutOfDateException,” on page 198
- ◆ Section A.3.34, “FactException,” on page 198
- ◆ Section A.3.35, “GridAuthenticationException,” on page 198
- ◆ Section A.3.36, “GridAuthorizationException,” on page 198
- ◆ Section A.3.37, “GridConfigurationException,” on page 198
- ◆ Section A.3.38, “GridDeploymentException,” on page 198
- ◆ Section A.3.39, “GridException,” on page 198
- ◆ Section A.3.40, “GridObjectNotFoundException,” on page 199

A.3.1 AgentListener

Provides the interface necessary for processing messages sent from the Orchestrate Server.

For complete documentation, see [AgentListener \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/AgentListener.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/AgentListener.html).

A.3.2 ClientAgent

API for client communication with server for job and datagrid operations.

For complete documentation, see [ClientAgent \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/ClientAgent.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/ClientAgent.html).

A.3.3 Credential

A credential used for identity on the PlateSpin Orchestrate system.

For complete documentation, see [Credential \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Credential.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Credential.html).

A.3.4 Fact

The Fact object.

For complete documentation, see [Fact](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Fact.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Fact.html).

A.3.5 FactSet

Definition of a set of facts.

For complete documentation, see [FactSet](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/FactSet.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/FactSet.html).

A.3.6 GridObjectInfo

Client interface to any Grid object.

For complete documentation, see [GridObjectInfo](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/GridObjectInfo.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/GridObjectInfo.html).

A.3.7 ID

A unique identifier for an engine or a facility or Grid object.

For complete documentation, see [ID](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/ID.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/ID.html).

A.3.8 JobInfo

A client representation of a deployed job.

For complete documentation, see [JobInfo](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/JobInfo.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/JobInfo.html).

A.3.9 Message

A base interface for all the messages in the system.

For complete documentation, see [Message](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.html).

A.3.10 Message.Ack

A general acknowledgement of “action” message.

For complete documentation, see [Message.Ack](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.Ack.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.Ack.html).

A.3.11 Message.AuthFailure

Authentication failure messages indicates that processing of a client message will not occur because client credentials are invalid.

For complete documentation, see [Message.AuthFailure](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.AuthFailure.html) (http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.AuthFailure.html).

A.3.12 Message.ClientResponseMessage

Message All messages that can optionally carry an error string back to the client extend this.

For complete documentation, see [Message.ClientResponseMessage \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.ClientResponseMessage.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.ClientResponseMessage.html).

A.3.13 Message.ConnectionID

All messages that can optionally carry an error string back to the client extend this.

For complete documentation, see [Message.ConnectionID \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.ConnectionID.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.ConnectionID.html).

A.3.14 Message.Event

An Event is used to signal clients and workflows.

For complete documentation, see [Message.Event \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.Event.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.Event.html).

A.3.15 Message.GetGridObjects

Client request to retrieve an (optionally ordered) set of grid objects that match a search criteria (constraint).

For complete documentation, see [Message.GetGridObjects \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.GetGridObjects.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.GetGridObjects.html).

A.3.16 Message.GridObjects

Server response to client request to retrieve grid a grid object set.

For complete documentation, see [Message.GridObjects \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.GridObjects.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.GridObjects.html).

A.3.17 Message.JobAccepted

A JobAccepted message is sent in response to a RunJob message when a job is successfully accepted into the system.

For complete documentation, see [Message.JobAccepted \(http://www.novell.com/documentation/zen_orchestrator13/reference/javadoc/com/novell/zos/grid/Message.JobAccepted.html\)](http://www.novell.com/documentation/zen_orchestrator13/reference/javadoc/com/novell/zos/grid/Message.JobAccepted.html).

A.3.18 Message.JobError

A JobError message is sent when an unrecoverable error occurs in a job.

For complete documentation, see [Message.JobError \(http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobError.html\)](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobError.html).

A.3.19 Message.JobFinished

A JobFinished message is sent when processing of a job completes.

For complete documentation, see [Message.JobFinished \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobFinished.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobFinished.html).

A.3.20 Message.JobIdEvent

Base Event interface for retrieving JobID used for jobid messages.

For complete documentation, see [Message.JobIdEvent \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobIdEvent.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobIdEvent.html).

A.3.21 Message.JobInfo

A JobInfo message contains information describing a deployed job.

For complete documentation, see [Message.JobInfo \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobInfo.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobInfo.html).

A.3.22 Message.Jobs

A Jobs message contains a list of deployed job names.

For complete documentation, see [Message.Jobs \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.Jobs.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.Jobs.html).

A.3.23 Message.JobStarted

A JobStarted message is sent when a job is successfully started.

For complete documentation, see [Message.JobStarted \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobStarted.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobStarted.html).

A.3.24 Message.JobStatus

A JobStatus message contains the state of the specified job.

For complete documentation, see [Message.JobStatus \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobStatus.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.JobStatus.html).

A.3.25 Message.LoginFailed

Response message for an unsuccessful login

For complete documentation, see [Message.LoginFailed \(http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.LoginFailed.html\)](http://www.novell.com/documentation/ps0_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.LoginFailed.html).

A.3.26 Message.LoginSuccess

Response message for a successful login.

For complete documentation, see [Message.LoginSuccess](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.LoginSuccess.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.LoginSuccess.html).

A.3.27 Message.LogoutAck

A LogoutAck indicates success or failure of logout operation.

For complete documentation, see [Message.LogoutAck](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.LogoutAck.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.LogoutAck.html).

A.3.28 Message.RunningJobs

A RunningJobs message contains the list of running jobs.

For complete documentation, see [Message.RunningJobs](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.RunningJobs.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.RunningJobs.html).

A.3.29 Message.ServerStatus

A ServerStatus message.

For complete documentation, see [Message.ServerStatus](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.ServerStatus.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Message.ServerStatus.html).

A.3.30 Node

Internal interface for Node (Resource) Grid object.

For complete documentation, see [Node](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Node.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Node.html).

A.3.31 Priority

Priority information.

For complete documentation, see [Priority](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Priority.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/Priority.html).

A.3.32 WorkflowInfo

A WorkflowInfo can represent either a snapshot of a running instance or an historical record of an instance.

For complete documentation, see [WorkflowInfo](http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/WorkflowInfo.html) (http://www.novell.com/documentation/psa_orchestrate20/resources/javadoc/com/novell/zos/grid/WorkflowInfo.html).

A.3.33 ClientOutOfDateException

Grid exception indicating the client is not compatible with the server.

For complete documentation, see [ClientOutOfDateException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/ClientOutOfDateException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/ClientOutOfDateException.html).

A.3.34 FactException

For exceptions that occur in accessing or setting facts.

For complete documentation, see [FactException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/FactException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/FactException.html).

A.3.35 GridAuthenticationException

Thrown when authentication is denied by a PlateSpin Orchestrate Server.

For complete documentation, see [GridAuthenticationException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridAuthenticationException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridAuthenticationException.html).

A.3.36 GridAuthorizationException

Thrown when credentials are insufficient for the desired grid operation.

For complete documentation, see [GridAuthorizationException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridAuthorizationException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridAuthorizationException.html).

A.3.37 GridConfigurationException

Grid exception thrown to indicate a Grid configuration error.

For complete documentation, see [GridConfigurationException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridConfigurationException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridConfigurationException.html).

A.3.38 GridDeploymentException

Thrown when credentials are insufficient for the desired grid operation.

For complete documentation, see [GridDeploymentException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridDeploymentException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridDeploymentException.html).

A.3.39 GridException

The base exception for all Grid exceptions.

For complete documentation, see [GridException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridException.html).

A.3.40 GridObjectNotFoundException

Thrown when a grid object lookup does not find the requested object.

For complete documentation, see [GridObjectNotFoundException](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridObjectNotFoundException.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/grid/GridObjectNotFoundException.html).

A.4 TLS Package

Interface and classes used for secure authentication to Orchestrate Server:

- ♦ [Section A.4.1, “TlsCallbacks,” on page 199](#)
- ♦ [Section A.4.2, “PemCertificate,” on page 199](#)
- ♦ [Section A.4.3, “TlsConfiguration,” on page 199](#)

A.4.1 TlsCallbacks

Callback interface for TLS certificate exceptions.

For complete documentation of the class, see [TlsCallback](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/tls/TlsCallbacks.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/tls/TlsCallbacks.html).

A.4.2 PemCertificate

PEM Certificate wrapper for X.509 certificates.

For complete documentation of the class, see [PemCertificate](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/tls/PemCertificate.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/tls/PemCertificate.html).

A.4.3 TlsConfiguration

TLS Configuration parameters for Orchestrate Clients.

For complete documentation of the class, see [TlsConfiguration](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/tls/TlsConfiguration.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/tls/TlsConfiguration.html).

A.5 Toolkit Package

The Client agent, Constraint, and Credentials factory patterns used by the PlateSpin Orchestrate Server:

- ♦ [Section A.5.1, “ClientAgentFactory,” on page 200](#)
- ♦ [Section A.5.2, “ConstraintFactory,” on page 200](#)
- ♦ [Section A.5.3, “CredentialFactory,” on page 200](#)

A.5.1 ClientAgentFactory

Factory pattern used to create new clients for connection to a PlateSpin Orchestrate Server.

For complete documentation of the class, see [ClientAgentFactory](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/toolkit/ClientAgentFactory.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/toolkit/ClientAgentFactory.html).

A.5.2 ConstraintFactory

Factory pattern used to create constraint objects which may be combined into larger constraint hierarchies for use in searches or other constraint based matching.

For complete documentation of the class, see [ConstraintFactory](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/toolkit/ConstraintFactory.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/toolkit/ConstraintFactory.html).

A.5.3 CredentialFactory

Factory pattern used to create a credential used for connection to a PlateSpin Orchestrate Server.

For complete documentation of the class, see [CredentialFactory](http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/toolkit/CredentialFactory.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/javadoc/com/novell/zos/toolkit/CredentialFactory.html).

PlateSpin Orchestrate Job Classes and JDL Syntax

B

- ◆ [Section B.1, “Job Class,” on page 201](#)
- ◆ [Section B.2, “Joblet Class,” on page 201](#)
- ◆ [Section B.3, “Utility Classes,” on page 201](#)
- ◆ [Section B.4, “Built-in JDL Functions and Variables,” on page 201](#)
- ◆ [Section B.5, “Job State Field Values,” on page 203](#)
- ◆ [Section B.6, “Repository Information String Values,” on page 204](#)
- ◆ [Section B.7, “Joblet State Values,” on page 204](#)
- ◆ [Section B.8, “Resource Information Values,” on page 205](#)
- ◆ [Section B.9, “JDL Class Definitions,” on page 205](#)

B.1 Job Class

To review the detailed JDL structure of the joblet class, see [Job \(page 226\)](#).

B.2 Joblet Class

To review the detailed JDL structure of the joblet class, see [Joblet \(page 228\)](#).

B.3 Utility Classes

The following are some of the main utility JDL classes you can use to customize your PlateSpin[®] Orchestrate jobs:

- ◆ [DataGrid \(page 216\)](#)
- ◆ [Exec \(page 219\)](#)
- ◆ [MatrixInfo \(page 235\)](#)
- ◆ [ResourceInfo \(page 244\)](#)
- ◆ [RunJobSpec \(page 245\)](#)
- ◆ [ScheduleSpec \(page 246\)](#)

B.4 Built-in JDL Functions and Variables

The information in this section defines the built-in PlateSpin Orchestrate JDL functions and variables.

- ◆ [Section B.4.1, “getMatrix\(\),” on page 202](#)
- ◆ [Section B.4.2, “system\(cmd\),” on page 202](#)
- ◆ [Section B.4.3, “Grid Object TYPE_* Variables,” on page 202](#)

- ♦ [Section B.4.4, “The `__agent__` Variable,” on page 202](#)
- ♦ [Section B.4.5, “The `__jobname__` Variable,” on page 202](#)
- ♦ [Section B.4.6, “The `__mode__` Variable,” on page 203](#)

B.4.1 `getMatrix()`

This function returns the matrix grid object. For more information, see [MatrixInfo \(page 235\)](#).

Purpose: The matrix object is used to retrieve other grid objects in the system.

B.4.2 `system(cmd)`

This executes a system command in a shell on the resource. The command is passed to the operating system’s default command interpreter. On Microsoft Windows systems this is `cmd.exe`, while on POSIX systems, this is `/bin/sh`. Stdout and stderr are directed to the job log. No access to stdin is provided.

Returns: Returns an exit code result of the command execution.

B.4.3 Grid Object `TYPE_*` Variables

The list of variables are constants for grid object type. For more information, see [MatrixInfo \(page 235\)](#).

Variable Names:

```
TYPE_USER
TYPE_JOB
TYPE_RESOURCE
TYPE_VMHOST
TYPE_REPOSITORY
TYPE_USERGROUP
TYPE_JOBGROUP
TYPE_RESOURCEGROUP
TYPE_REPOSITORYGROUP
```

Type: String.

Purpose: Use these in JDL functions for retrieving and creating grid objects.

B.4.4 The `__agent__` Variable

Variable Name: `__agent__`

Type: Boolean.

Purpose: Defines whether the JDL is executing on the agent.

B.4.5 The `__jobname__` Variable

Variable Name: `__jobname__`

Type: String.

Purpose: Defines the name of the deployed job.

B.4.6 The `__mode__` Variable

Variable Name: `__mode__`

Type: String.

Purpose: Defines the execution mode.

Values:

`parse` - JDL is being parsed.

`deploy` - JDL is being deployed.

`undeploy` - JDL is being undeployed.

`runtime` - JDL is being executed.

B.5 Job State Field Values

Here are the job state field values for the [Job \(page 226\)](#) class:

Constant	Value	Description
<code>int CANCELLED_STATE</code>	9	Cancelled end state.
<code>int CANCELLING_STATE</code>	6	Cancelling. Transitions to: Cancelled.
<code>int COMPLETED_STATE</code>	8	Completed end state.
<code>int COMPLETING_STATE</code>	5	Completing. Transitions to: Completing.
<code>int FAILED_STATE</code>	10	Failed end state.
<code>int FAILING_STATE</code>	7	Failing. Transitions to: Failed.
<code>int PAUSED_STATE</code>	4	Paused. Transitions to: Running/Completing/ Failing/Cancelling.
<code>int QUEUED_STATE</code>	1	Queued. Transitions to: Starting/Failing/ Cancelling.
<code>int RUNNING_STATE</code>	3	Running. Transitions to: Paused/Completing/ Failing/Cancelling.
<code>int STARTING_STATE</code>	2	Starting. Transitions to: Running/Failing/ Cancelling.
<code>int SUBMITTED_STATE</code>	0	Submitted. Transitions to: Queued/Failing.
<code>String TERMINATION_TYPE_ADMIN</code>	"Admin"	Indicates Job was cancelled by the admin and only applies if Job is in <code>CANCELLED_STATE</code> . Value is obtained from <code>jobinstance.terminationtype fact.</code>

Constant	Value	Description
String TERMINATION_TYPE_JOB	“Job”	Indicates Job was cancelled due to exceeding the job timeout value and only applies if Job is in CANCELLED_STATE. The value is obtained from <code>jobinstance.terminationtype</code> fact.
String TERMINATION_TYPE_TIMEOUT	“Timeout”	Indicates Job was cancelled due to exceeding the job timeout value and only applies if Job is in CANCELLED_STATE. Value is obtained from <code>jobinstance.terminationtype</code> fact.
String TERMINATION_TYPE_USER	“User”	Indicate Job was cancelled by client user and only applies if Job is in CANCELLED_STATE. The value is obtained from <code>jobinstance.terminationtype</code> fact.

B.6 Repository Information String Values

Constant	Value	Description
SAN_TYPE_FibreChannel	Fibre Channel	Specifies a fibre channel SAN repository.
SAN_TYPE_ISCSI	iSCSI	Specifies an iSCSI SAN repository.
SAN_VENDOR_IQN	iqn	Specifies an IQN SAN repository.
SAN_VENDOR_NPIV	npiv	Specifies a N_Port ID Virtualization SAN repository.
TYPE_DATAGRID	datagrid	Specifies a datagrid repository.
TYPE_LOCAL	local	Specifies a local repository.
TYPE_NAS	NAS	Specifies a NAS repository.
TYPE_SAN	SAN	Specifies a SAN repository.
TYPE_VIRTUAL	virtual	Specifies a virtual repository.
TYPE_WAREHOUSE	warehouse	Specifies a warehouse repository.

B.7 Joblet State Values

The following values are defined for the various states that the joblet can be in:

Constant	Value	Description
INITIAL_STATE	0	Joblet initial state.
WAITING_STATE	1	Joblet waiting for a resource
WAITING_RETRY_STATE	2	Joblet waiting for a resource for retry.
CONTRACTED_STATE	3	Joblet waiting for a resource for retry.
STARTED_STATE	4	Joblet started on a resource.

Constant	Value	Description
PRE_CANCEL_STATE	5	Joblet starting cancellation.
CANCELLING_STATE	6	Joblet cancelling.
POST_CANCEL_STATE	7	Joblet finishing cancellation.
COMPLETING_STATE	8	Joblet completing state.
FAILING_STATE	9	Joblet failing state.
FAILED_STATE	11	Joblet failed end state.
CANCELLED_STATE	12	Joblet cancelled end state.
COMPLETED_STATE	13	Joblet completed end state.

See [Joblet \(page 228\)](#).

B.8 Resource Information Values

Use the following values to specify the

Constant	Value Type	Resource Description
TYPE_BM_INSTANCE	String	Blade server.
TYPE_BM_TEMPLATE	String	Blade server template.
TYPE_FIXED_PHYSICAL	String	Fixed physical server.
TYPE_VM_INSTANCE	String	VM server.
TYPE_VM_TEMPLATE	String	VM template.

For full class descriptions, see [ResourceInfo \(page 244\)](#).

B.9 JDL Class Definitions

The following PlateSpin Orchestrate JDL classes can be implemented in the custom jobs that you create. Because JDL is implemented in Java, we have provided direct links to detailed Javadoc for each of the “Pythonized” JDL classes below:

- ♦ [“AndConstraint\(\)” on page 207](#)
- ♦ [“BinaryConstraint” on page 208](#)
- ♦ [“BuildSpec” on page 209](#)
- ♦ [“CharRange” on page 210](#)
- ♦ [“ComputedFact” on page 211](#)
- ♦ [“ComputedFactContext” on page 212](#)
- ♦ [“Constraint” on page 213](#)
- ♦ [“ContainerConstraint” on page 214](#)

- ◆ “ContainsConstraint” on page 215
- ◆ “DataGrid” on page 216
- ◆ “DefinedConstraint” on page 217
- ◆ “EqConstraint” on page 218
- ◆ “Exec” on page 219
- ◆ “ExecError” on page 220
- ◆ “FileRange” on page 221
- ◆ “GeConstraint” on page 222
- ◆ “GridObjectInfo” on page 223
- ◆ “GroupInfo” on page 224
- ◆ “GtConstraint” on page 225
- ◆ “Job” on page 226
- ◆ “JobInfo” on page 227
- ◆ “Joblet” on page 228
- ◆ “JobletInfo” on page 229
- ◆ “JobletParameterSpace” on page 230
- ◆ “LeConstraint” on page 231
- ◆ “LtConstraint” on page 232
- ◆ “MatchContext” on page 233
- ◆ “MatchResult” on page 234
- ◆ “MatrixInfo” on page 235
- ◆ “MigrateSpec” on page 236
- ◆ “NeConstraint” on page 237
- ◆ “NotConstraint” on page 238
- ◆ “OrConstraint” on page 239
- ◆ “ParameterSpace” on page 240
- ◆ “PolicyInfo” on page 241
- ◆ “ProvisionSpec” on page 242
- ◆ “RepositoryInfo” on page 243
- ◆ “ResourceInfo” on page 244
- ◆ “RunJobSpec” on page 245
- ◆ “ScheduleSpec” on page 246
- ◆ “Timer” on page 247
- ◆ “UndefinedConstraint” on page 248
- ◆ “UserInfo” on page 249
- ◆ “VMHostInfo” on page 250
- ◆ “VmSpec” on page 251

AndConstraint()

Representation of the And Constraint. Perform a logical ANDing of all child constraints. If this constraint contains no children, no operation is performed. Constraints are added to this constraint using `add()`.

See Also

- ◆ [ContainerConstraint \(page 214\)](#)
- ◆ Javadoc: *AndConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/AndConstraint.html)

BinaryConstraint

Representation of a Constraint operating on the left and right operands. This is a base class and is not directly constructed.

See Also

- ◆ Subclasses: [ContainsConstraint \(page 215\)](#), [EqConstraint \(page 218\)](#), [GeConstraint \(page 222\)](#), [GtConstraint \(page 225\)](#), [LeConstraint \(page 231\)](#), [LtConstraint \(page 232\)](#), [NeConstraint \(page 237\)](#).
- ◆ Javadoc: [*BinaryConstraint*](http://www.novell.com/documentation/pso_orchestrate20/resources/jdldoc/com/novell/zos/jdl/BinaryConstraint.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/jdldoc/com/novell/zos/jdl/BinaryConstraint.html)

BuildSpec

Defines the attributes for building a new VM. An instance of this class is passed to `resource.build()`.

See Also

- ♦ Javadoc: *BuildSpec* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/BuildSpec.html)

CharRange

Define lexical character string range of values for ParameterSpace scheduling.

See Also

- ♦ Javadoc: *CharRange* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/CharRange.html)

ComputedFact

Defines the base class for creating custom computed facts. Computed facts provide the ability to create custom calculations that extend the built-in factsets for a grid object. The computed fact can be in constraints. User defined computed facts are required to subclass this class. In order to use `ComputedFact`, you must deploy a subclass of `ComputedFact` and then create a linked fact referencing the deployed `ComputedFact`. The linked fact is then used in constraints.

See Also

- ♦ [ComputedFactContext \(page 212\)](#)
- ♦ Javadoc: *ComputedFact* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/ComputedFact.html)

ComputedFactContext

Provides access to the evaluation context. See Example below.

Description

The context contains the grid objects that the constraint engine uses to evaluate constraints. If they are available in the current context, the `ComputedFactContext` provides access to the current job instance, deployed job, user, resource, VM host, and repository grid objects.

The VM host and repository grid objects are only in the context for the evaluation of the provisioning constraints such as `vmHost`. The job and job instance objects are only in the context for a resource or allocation constraint evaluation.

See Also

- ♦ [ComputedFact \(page 211\)](#)
- ♦ Javadoc: [ComputedFactContext](http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/ComputedFactContext.html) (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/ComputedFactContext.html)

Constraint

Defines the base class for all constraint classes.

See Also

- ♦ [BinaryConstraint \(page 208\)](#), [ContainerConstraint \(page 214\)](#), [DefinedConstraint \(page 217\)](#), [UndefinedConstraint \(page 248\)](#).
- ♦ Javadoc: *Constraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/Constraint.html)

ContainerConstraint

Representation of a Constraint that contains other Constraints. This is a base class and is not directly constructed.

See Also

- ◆ Subclasses: [AndConstraint\(\)](#) (page 207), [NotConstraint](#) (page 238), [OrConstraint](#) (page 239)
- ◆ Javadoc: *ContainerConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdldoc/com/novell/zos/jdl/ContainerConstraint.html)

ContainsConstraint

Representation of the Contains Constraint. Evaluates to true only if the left side fact is defined in the match context. If the left side is not defined, this will evaluate to False. Contains is typically used to check membership of a value in a group fact.

See Also

Javadoc: *ContainsConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/ContainsConstraint.html)

DataGrid

General interface to the datagrid. See [Chapter 3, “The PlateSpin Orchestrate Datagrid,”](#) on page 25.

See Also

- ♦ [GridObjectInfo \(page 223\)](#)
- ♦ Javadoc: *DataGrid* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/DataGrid.html)

DefinedConstraint

Representation of the Defined Constraint. Evaluates to true only if the left side fact is defined in the match context. If the left side is not defined, this will evaluate to False. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints.

See Also

- ♦ [Constraint \(page 213\)](#), [ContainerConstraint \(page 214\)](#), and [ContainsConstraint \(page 215\)](#)
- ♦ Javadoc: *DefinedConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/DefinedConstraint.html)

EqConstraint

Representation of the Equals Constraint. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints. Extends [BinaryConstraint \(page 208\)](#).

See Also

- ♦ [BinaryConstraint \(page 208\)](#)
- ♦ Javadoc: *EqConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/EqConstraint.html)

Exec

The Exec class is used to manage command line execution on resources. This class defines options for input, output and error stream handling, and process management including signaling, error and timeout control.

Description

A command's standard output and error can be redirected to a file, to a stream, to write to the job log, or be discarded. By default, the output is discarded. A command's standard input can be directed from a file or a stream can be written to. By default, the input is not used.

By default, command line execution is done in behalf of the job user. Exec instances are only allowed during the running of the Joblet class on a resource. The built-in function `system()` can also be used for simple execution of command lines.

See Also

- ♦ [BinaryConstraint \(page 208\)](#) and [ExecError \(page 220\)](#)
- ♦ Javadoc: *Exec* (http://www.novell.com/documentation/psoc_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/Exec.html)

ExecError

`ExecError` is raised for errors in executing a command line using the [Exec \(page 219\)](#) class or `system()`. Normal raising of this error causes the joblet to fail. Put this Error in an try except block to handle the error.

See Also

- ◆ [Exec \(page 219\)](#)
- ◆ Javadoc: *ExecError* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/ExecError.html)

FileRange

Define a range of values for a [ParameterSpace \(page 240\)](#) based on the lines of a text file. An instance of this class is used as a dimension in a `ParameterSpace` definition. The file name must either refer to a file that is readable from the server and resources (on a shared file system) or must be a [DataGrid \(page 216\)](#) file URL.

See Also

- ♦ [DataGrid \(page 216\)](#) and [ParameterSpace \(page 240\)](#)
- ♦ Javadoc: *FileRange* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/FileRange.html)

GeConstraint

Representation of the Greater than or Equals constraint. Performs a ‘greater than or equal to’ constraint operation. Missing arguments will always result in this constraint evaluating to false. The standard lexicographical ordering of values is used to determine result. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints. Extends [BinaryConstraint \(page 208\)](#).

See Also

- ♦ [Constraint \(page 213\)](#) and [BinaryConstraint \(page 208\)](#).
- ♦ Javadoc: *GeConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/GeConstraint.html)

GridObjectInfo

The `GridObjectInfo` class is the base class representation of all grid objects in the system. This provides accessors and setters to a grid object's fact set.

See Also

Javadoc: *GridObjectInfo* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/GridObjectInfo.html)

GroupInfo

The GroupInfo class is a representation of Group grid objects. Operations include retrieving the group member lists and adding/removing from the group member lists, and retrieving and setting facts on the group. Extends [GridObjectInfo \(page 223\)](#).

See Also

- ◆ [GridObjectInfo \(page 223\)](#)
- ◆ Javadoc: *GroupInfo* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/GroupInfo.html)

GtConstraint

Representation of the Greater than Constraint. Performs a 'greater than' constraint operation. Missing arguments will always result in this constraint evaluating to false. The standard lexicographical ordering of values is used to determine result. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints.

See Also

- ♦ [Constraint \(page 213\)](#)
- ♦ Javadoc: *GtConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/GtConstraint.html)

Job

The `Job` class represents a running job instance. This class defines functions for interacting with the server including handling notification of job state transitions, child job submission, managing joblets and for receiving and sending events from resources and from clients. A job writer defines a subclass of the `Job` class and uses the methods available on the `Job` class for scheduling joblets and event processing.

See Also

- ♦ [JobInfo \(page 227\)](#), [Joblet \(page 228\)](#), [JobletInfo \(page 229\)](#), [JobletParameterSpace \(page 230\)](#)
- ♦ Javadoc: *Job* (http://www.novell.com/documentation/ps0_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/Job.html)

JobInfo

The JobInfo class is a representation of a deployed job. The factset available on the JobInfo class is the aggregation of the job's policy and policies on the groups the job is a member of. This includes the "job.*" and "jobargs.*" fact namespaces.

See Also

- ♦ [Job \(page 226\)](#)
- ♦ Javadoc: *JobInfo* (http://www.novell.com/documentation/psa_orchestrate20/resources/jdldoc/com/novell/zos/jdl/JobInfo.html)

Joblet

Defines the attributes for creating a virtual machine. An instance of this class is passed to `resource.createInstance()`, `resource.createTemplate()`, `resource.clone()`.

See Also

- ◆ Job, JobInfo
- ◆ Javadoc: *Joblet* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/Joblet.html)

JobletInfo

JobletInfo is a representation of the joblet grid object created when a job calls `schedule()` to create joblets. This class provides access to a joblet's factset and operations on a joblet such as cancellation and sending events to a joblet that is running on a resource. The separate Joblet class defines execution on a resource.

See Also

Javadoc: *JobletInfo* (http://www.novell.com/documentation/psa_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/JobletInfo.html)

JobletParameterSpace

JobletParameterSpace is a slice of the ParameterSpace allocated to a joblet. As the scheduler defines slices of the parameter space for a given schedule(), JobletParameterSpace instances are created for each joblet. This slice of the parameter space is delivered to the resource on joblet execution. The JobletParameterSpace can also be retrieved from the joblet object.

See Also

Javadoc: *JobletParameterSpace* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/JobletParameterSpace.html)

LeConstraint

Representation of the Less than or equals Constraint. Performs a 'less than or equal to' constraint operation. Missing arguments will always result in this constraint evaluating to false. The standard lexicographical ordering of values is used to determine result. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints.

See Also

Javadoc: *LeConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdldoc/com/novell/zos/jdl/LeConstraint.html)

LtConstraint

Representation of the Less than Constraint. Performs a "less than" constraint operation. Missing arguments always result in this constraint evaluating to false. The standard lexicographical ordering of values is used to determine result. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints.

See Also

Javadoc: *LtConstraint* (http://www.novell.com/documentation/ps0_orchestrate20/resources/jdldoc/com/novell/zos/jdl/LtConstraint.html)

MatchContext

The MatchContext defines a context for evaluating a constraint. An instance of this class is supplied to `match()` for evaluating constraints. The `MatchContext` provides a way to setup an evaluation context that the constraint engine is using to evaluate constraints. The `MatchContext` is filled out with the context that is required for evaluating your constraints. This includes assigning a deployed Job, User, Resource, VM Host and Repository Grid objects.

See Also

- ♦ Javadoc: *MatchContext* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdldoc/com/novell/zos/jdl/MatchContext.html)

MatchResult

The `MatchResult` class defines results of a `Constraint` `match()`. Instance of `MatchResult` is returned from `match()` operations. From the `MatchResult` you can retrieve a list of the IDs of the matching Grid objects and the non-matching Grid objects.

See Also

- ♦ Javadoc: *MatchResult* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/MatchResult.html)

MatrixInfo

The `MatrixInfo` class is a representation of the matrix grid object (see [GridObjectInfo \(page 223\)](#)). This provides operations for retrieving and creating grid objects in the system. `MatrixInfo` is retrieved using the built-in `getMatrix()` function. Write capability is dependent on the context in which `getMatrix()` is called. For example, in a joblet process on a resource, creating new grid objects is not supported.

See Also

- ♦ Javadoc: *MatrixInfo* (http://www.novell.com/documentation/psa_orchestrate20/resources/jddoc/com/novell/zos/jdl/MatrixInfo.html)
- ♦ [Section B.4, “Built-in JDL Functions and Variables,” on page 201.](#)

MigrateSpec

The MigrateSpec class defines the options for the migrate action. An instance of this class is passed to the `resource.migrate()` method.

Example

The following is an example of using MigrateSpec to define a migrate action for a Virtual Machine named "sles10" to a VM Host named "host2:"

```
vm = getMatrix().getGridObject(TYPE_RESOURCE, "sles10")
spec = MigrateSpec()
spec.setHost('host2')
vm.migrate(spec)
```

See Also

Javadoc: *MigrateSpec* (http://www.novell.com/documentation/ps0_orchestrate20/resources/jdldoc/com/novell/zos/jdl/MigrateSpec.html)

NeConstraint

Representation of the Not Equals Constraint. Performs a not equal constraint operation. Missing arguments will always result in this constraint evaluating to false. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints.

See Also

Javadoc: *NeConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/NeConstraint.html)

NotConstraint

Representation of a Not Constraint Object. Performs a logical not operation of all the child constraints. This is a no-op if this constraint contains no children. Constraints are added to this constraint using `add()`.

See Also

- ◆ See [Constraint \(page 213\)](#) and [ContainerConstraint \(page 214\)](#).
- ◆ Javadoc: *NotConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/NotConstraint.html)

OrConstraint

Representation of Or Constraint Object. Perform a logical or-ing operation of all the child constraints. This is a no-op if this constraint contains no children. Constraints are added to this constraint using `add()`.

See Also

- ◆ See [Constraint \(page 213\)](#) and [ContainerConstraint \(page 214\)](#).
- ◆ Javadoc: *OrConstraint* (http://www.novell.com/documentation/psa_orchestrate20/resources/jldoc/com/novell/zos/jdl/OrConstraint.html)

ParameterSpace

Defines a parameter space to be used by the scheduler to create a joblet set. A parameter space may consist of rows of columns or a list of columns that is expanded and can be turned into a cross product. Use `appendRow` to create a `rowMajor` parameter space or `appendCol` to define a column expansion. You cannot use both `appendRow()` and `appendCol()` in the same `ParameterSpace`. Once the scheduler defines a slice of the parameter space for a given joblet, the scheduler creates `JobletParameterSpace` instances for each joblet. This slice of the parameter space is delivered to the resource.

See Also

Javadoc: *ParameterSpace* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/ParameterSpace.html)

PolicyInfo

Representation of a Policy Object. This class allows for associating and unassociation of Grid objects using this policy

See Also

Javadoc: *PolicyInfo* (http://www.novell.com/documentation/psf_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/PolicyInfo.html)

ProvisionSpec

Defines the attributes for starting a provision. An instance of this class is passed to `self.provision()`. Defining a provision to reserve a provisioned resource "nightly" for a user is an instance of when this function might be used: `spec = ProvisionSpec()`
`spec.setReserveForUser('nightly') self.provision(spec)`

See Also

Javadoc: *ProvisionSpec* (http://www.novell.com/documentation/pso_orchestrate20/resources/jldoc/com/novell/zos/jdl/ProvisionSpec.html)

RepositoryInfo

`RepositoryInfo` is a representation of a repository grid object. This class provides accessors and setters for Repository facts. See [MatrixInfo \(page 235\)](#) for how to script creation of Repository objects.

See Also

- ◆ See [GridObjectInfo \(page 223\)](#) and [MatrixInfo \(page 235\)](#).
- ◆ Javadoc: *RepositoryInfo* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/RepositoryInfo.html)

ResourceInfo

ResourceInfo is a representation of a resource grid object. This class inherits the base fact operations from [GridObjectInfo \(page 223\)](#) and adds the provisioning operations for provisionable resources such as VMs. See [MatrixInfo \(page 235\)](#) for how to script creation of Resource objects.

See Also

- ◆ [GridObjectInfo \(page 223\)](#) and [MatrixInfo \(page 235\)](#).
- ◆ Javadoc: *ResourceInfo* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/ResourceInfo.html)

RunJobSpec

Defines the attributes for starting a child job. An instance of this class is passed to `self.runJob()`.

See Also

Javadoc: *RunJobSpec* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdldoc/com/novell/zos/jdl/RunJobSpec.html)

ScheduleSpec

Defines one or more joblets to be scheduled and run on resources. A `ScheduleSpec` instance is passed to the job's `schedule()`. `schedule()` creates the joblets and schedules joblets to run on resources.

See Also

- ◆ [Joblet \(page 228\)](#)
- ◆ Javadoc: *ScheduleSpec* (http://www.novell.com/documentation/ps0_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/ScheduleSpec.html)

Timer

Timer schedules a callback to a job or joblet method. Timers can schedule a one time or a repeated callback on an interval basis. Any Timers created in a job or joblet are shut down on job or joblet completion.

See Also

Javadoc: *Timer* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/Timer.html)

UndefinedConstraint

Representation of the Undefined Constraint. Evaluates to true only if the left side fact is *not* defined in the match context. If the left side is not defined, this will evaluate to `false`. This constraint can be used independently or added to a And, Or, Not constraint to combine with other constraints.

See Also

- ♦ [Constraint \(page 213\)](#)
- ♦ Javadoc: *UndefinedConstraint* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/UndefinedConstraint.html)

UserInfo

UserInfo is a representation of a user grid object. This class provides accessors and setters for User facts.

See Also

Javadoc: *UserInfo* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/UserInfo.html)

VMHostInfo

The `VMHostInfo` class is a representation of a virtual machine host grid object. This class provides accessors and setters to the VM host facts and operations to control the state of the VM host object.

See Also

- ♦ [VmSpec \(page 251\)](#)
- ♦ Javadoc: *VMHostInfo* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/VMHostInfo.html)

VmSpec

Defines the attributes for creating a virtual machine. An instance of this class is passed to `resource.createInstance()`, `resource.createTemplate()`, `resource.clone()`.

Example

Example of using `VmSpec` for creating a clone on a named host from a template resource:

```
template = getMatrix().getGridObject(TYPE_RESOURCE, "myTemplate")
spec = VmSpec()
spec.setNewName("newvm")
spec.setHost('vmhost-qa')
template.clone(spec)
```

If the host and repository is not set, the default is to use the source resource object's repository as the destination repository.

See Also

- ♦ [VMHostInfo \(page 250\)](#)
- ♦ Javadoc: *VmSpec* (http://www.novell.com/documentation/pso_orchestrate20/resources/jdl/doc/com/novell/zos/jdl/VmSpec.html)

Understanding Resource Metrics Facts

C

When you install the PlateSpin® Orchestrate Agent on a machine, you can optionally install the Orchestrate Monitoring Agent along with it. The Monitoring Agent uses the Ganglia Monitoring Daemon (gmond) to automatically collect metrics and send them to the Orchestrate Monitoring Server. You can use the following command to check the status of an installed Monitoring Agent:

```
# /etc/init.d/novell-gmond status
```

If the daemon is operating normally, it returns a `running` status.

When you install and configure the Orchestrate Monitoring Agent (gmond), it is set by default to report metrics on port 8649, which is also detected by the Orchestrate Agent. When communication is established, the gmond daemon sends out metrics data, which are then gathered by the Orchestrate Agent and set as fact values associated with the resource where the daemon is running. You can verify the connection with the following command:

```
telnet localhost 8649
```

If gmond is running and communicating properly, an XML document listing the reported metrics is displayed.

This section includes information about the resource metrics facts that are gathered, the unit conversion performed by Orchestrate on the Ganglia-provided values, and how you can use these facts to help you manage the resources in the grid.

- ♦ [Section C.1, “Resource Facts,” on page 253](#)
- ♦ [Section C.2, “Interpreting the Units of Metrics Fact Values,” on page 254](#)

C.1 Resource Facts

The Orchestrate Agent uses the metrics collected by gmond to create fact values for a given resource. These facts are therefore externally generated and are not among the default facts reported by the PlateSpin Orchestrate Agent. The agent updates these externally generated fact values every 30 seconds. All of these fact values have a `resource.metrics.` prefix.

For example, gmond collects a metrics value called `load_one`. The Orchestrate Agent sets this value as the `resource.metrics.load_one` fact.

To see a list of these facts in the Orchestrate Development Client,

- 1 In the Explorer panel, select a resource.
- 2 In the Workspace panel, select *Constraints/Facts*.

The names of the resource metrics facts are displayed in bold font because they were added as new facts to the default fact list. The following sample is a list of the default Ganglia-generated metrics facts with data type and an example value:

```

<fact name="resource.metrics.boottime" value="1239122234.0000" type="Real" />
<fact name="resource.metrics.bytes_in" value="208.8800" type="Real" />
<fact name="resource.metrics.bytes_out" value="68.9700" type="Real" />
<fact name="resource.metrics.cpu_idle" value="76.9000" type="Real" />
<fact name="resource.metrics.cpu_nice" value="95.2000" type="Real" />
<fact name="resource.metrics.cpu_num" value="0.0000" type="Real" />
<fact name="resource.metrics.cpu_speed" value="2" type="Integer" />
<fact name="resource.metrics.cpu_system" value="1596" type="Integer" />
<fact name="resource.metrics.cpu_user" value="0.3000" type="Real" />
<fact name="resource.metrics.cpu_wio" value="4.0000" type="Real" />
<fact name="resource.metrics.disk_free" value="0.4000" type="Real" />
<fact name="resource.metrics.disk_total" value="27090" type="Integer" />
<fact name="resource.metrics.gexec" value="48213" type="Integer" />
<fact name="resource.metrics.gexec" value="OFF" type="String" />
<fact name="resource.metrics.load_fifteen" value="0.2000" type="Real" />
<fact name="resource.metrics.load_five" value="0.4100" type="Real" />
<fact name="resource.metrics.load_one" value="1.1900" type="Real" />
<fact name="resource.metrics.machine_type" value="x86" type="String" />
<fact name="resource.metrics.mem_buffers" value="299" type="Integer" />
<fact name="resource.metrics.mem_cached" value="761" type="Integer" />
<fact name="resource.metrics.mem_free" value="65" type="Integer" />
<fact name="resource.metrics.mem_shared" value="0" type="Integer" />
<fact name="resource.metrics.mem_total" value="1989" type="Integer" />
<fact name="resource.metrics.os_name" value="Linux" type="String" />
<fact name="resource.metrics.os_release" value="2.6.27.19-5-pae" type="String" />
/>
<fact name="resource.metrics.part_max_used" value="70.8000" type="Real" />
<fact name="resource.metrics.part_max_used.units" value="" type="String" />
<fact name="resource.metrics.pkts_in" value="0.4500" type="Real" />
<fact name="resource.metrics.pkts_out" value="0.6300" type="Real" />
<fact name="resource.metrics.proc_run" value="0" type="Integer" />
<fact name="resource.metrics.proc_total" value="411" type="Integer" />
<fact name="resource.metrics.swap_free" value="2039" type="Integer" />
<fact name="resource.metrics.swap_total" value="2047" type="Integer" />
<fact name="resource.metrics.vm_type" value="" type="String" />
<fact name="resource.metrics.vm_type.units" value="" type="String" />

```

These are the metrics reported in Orchestrate systems that use the `gmond.conf` created when Orchestrate Monitoring Agent was installed and configured. The open source `gmond` might include other metrics that can be monitored. You can modify the default Orchestrate `gmond` configuration file to report these metrics after it is initially installed and configured. For information about modifying the file, see the [gmond.conf man page \(http://linux.die.net/man/5/gmond.conf\)](http://linux.die.net/man/5/gmond.conf).

By using the XML constraint language, you can utilize these resource metrics facts as you would use any other fact in PlateSpin Orchestrate. For example, you could create an Event that sets thresholds for the amount of incoming network packets. When that threshold is exceeded, a Scheduled Job could be triggered or a notification e-mail sent. For more information, see [Section 7.14, “Using an Event Notification in a Job,” on page 66](#).

C.2 Interpreting the Units of Metrics Fact Values

The Orchestrate Agent converts most of the Ganglia metrics values to PlateSpin Orchestrate standard units. This allows fact values to be compared in constraints without the need to perform conversions explicitly. In cases where units are not known or cannot be converted, a separate fact with a `.units` suffix is included. For example:

```

<fact name="resource.metrics.bytes_in" value="bytes/sec" type="String" />

```

The following table lists the `resource.metrics` facts and the units of measure used for each fact value:

Table C-1 *Resource*

Resource Metric Fact With Reported Value	Orchestrate Measurement Unit of the Value
boottime	32-bit seconds timestamp
bytes_in	byte rate measured in bytes per second
bytes_out	byte rate measured in bytes per second
cpu_idle	percentage
cpu_idle	percentage
cpu_nice	percentage
cpu_num	number of CPUs
cpu_speed	megahertz as an integer
cpu_system	percentage
cpu_user	percentage
cpu_wio	percentage
disk_total	total in binary megabytes
disk_free	total in binary megabytes
gexec	simple string
load_fifteen	real number
load_five	real number
load_one	real number
machine_type	simple string
mem_buffers	memory in megabytes (integer)
mem_cached	memory in megabytes (integer)
mem_free	memory in megabytes (integer)
mem_shared	memory in megabytes (integer)
mem_total	memory in megabytes (integer)
os_name	simple string
os_release	simple string
pkts_in	packet rate in packets per second
pkts_out	packet rate in packets per second
proc_run	processes run (integer)
proc_total	total processes (integer)

Resource Metric Fact With Reported Value	Orchestrate Measurement Unit of the Value
swap_free	memory in megabytes (integer)
swap_total	memory in megabytes (integer)

Documentation Updates

D

This section contains information about documentation content changes that were made in this *PlateSpin Orchestrate Developer Guide and Reference* after the initial release of PlateSpin Orchestrate 2.0. 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, “July 9, 2009,” on page 257](#)
- ♦ [Section D.2, “June 17, 2009 \(2.0.2 Release\),” on page 257](#)

D.1 July 9, 2009

Updates were made to the following sections:

Location	Update
Section 9.5, “Provisioning Virtual Machines,” on page 82	Moved two items to the end of the (first) bullet list, then indicated in a note there that VMware Server and VMware ESX Hypervisors employ a provisioning technology considered to be “experimental” by Novell. Added a similar note following Figure 9-2 on page 84 .

D.2 June 17, 2009 (2.0.2 Release)

Updates were made to the following sections:

Location	Update
Appendix B, “PlateSpin Orchestrate Job Classes and JDL Syntax,” on page 201	Added two new classes, MatchContext (page 233) and MatchResult (page 234) .
Appendix C, “Understanding Resource Metrics Facts,” on page 253	New information in a new appendix.

