

Job Manager Guide

Novell® PlateSpin® Orchestrate

2.0

February 23, 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 © 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	7
1 The Role of the Job Manager	9
2 Using the PlateSpin Orchestrate User Portal	11
2.1 Getting Started with the Orchestrate User Portal	11
2.2 Using the Orchestrate User Portal	11
2.2.1 Connecting to the Orchestrate User Portal	11
2.2.2 Displaying Available Jobs	13
2.2.3 Running a Job	14
2.2.4 Monitoring a Job	19
2.2.5 Viewing the Job Log	20
2.2.6 Pausing and Resuming a Running Job	20
2.2.7 Canceling a Running Job.	21
2.2.8 Logging Out of the Orchestrate User Portal	21

About This Guide

This *Job Manager Guide* introduces the role and the User Portal tool used by a PlateSpin® Orchestrate job manager, also referred to as the PlateSpin Orchestrate “end user.” The guide provides an overview of the web-based Orchestrate User Portal and explains how to use the portal to monitor and manage applications running on the PlateSpin Orchestrate Server. You can use this reference along with the information in “[The zos Command Line Tool](#)” section in the *PlateSpin Orchestrate 2.0 Command Line Reference*. Along with the PlateSpin Orchestrate User Portal, you can use the zos command line tool to monitor jobs. The guide is organized as follows:

- ♦ [Chapter 1, “The Role of the Job Manager,” on page 9](#)
- ♦ [Chapter 2, “Using the PlateSpin Orchestrate User Portal,” on page 11](#)

Audience

The contents of this guide are of interest to the end user of PlateSpin Orchestrate, also called a “job manager,” runs and manages jobs that have been created by a job developer and deployed by the administrator. It is also possible that the end user could be a developer who has created applications to run on distributed computing resources.

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.

Feedback

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

Additional Product Documentation

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

- ♦ [PlateSpin Orchestrate 2.0 Getting Started Reference](#)
- ♦ [PlateSpin Orchestrate 2.0 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 Developer Guide and Reference*

Documentation Updates

This *Job Manager Guide* is updated periodically. To view the most recent version, visit the [PlateSpin Orchestrate 2.0 documentation Web site \(http://www.novell.com/documentation/ps_orchestrate20/\)](http://www.novell.com/documentation/ps_orchestrate20/) and look for the documentation updates section included at the end of the book.

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.

The Role of the Job Manager

1

The Job Manager has permission to run jobs on the PlateSpin® Orchestrate Server. These jobs are deployed by the PlateSpin Orchestrate administrator.

Before you begin using the User Portal in your role as a job manager , you must have:

- ♦ A User account created by the PlateSpin Orchestrate Administrator.
- ♦ An Internet browser.

The following browsers support the User Portal:

- ♦ Internet Explorer* version 6.0 or higher
- ♦ Netscape* Navigator* version 6.0 or higher
- ♦ Mozilla* version 1.5 or higher
- ♦ Firefox* 2.0 or higher

Using the PlateSpin Orchestrate User Portal

2

The PlateSpin® Orchestrate User Portal is a web-based management tool that makes it easy to use the PlateSpin Orchestrate Server to monitor and run jobs. This section provides an introduction and overview of the portal and includes the following information:

- ♦ [Section 2.1, “Getting Started with the Orchestrate User Portal,” on page 11](#)
- ♦ [Section 2.2, “Using the Orchestrate User Portal,” on page 11](#)

2.1 Getting Started with the Orchestrate User Portal

A network or grid management administrator provides users with a URL (the address of the Orchestrate Server) to log in to the portal application.

Before using the portal, users must specify a username and password created by a PlateSpin Orchestrate administrator to log in to the Orchestrate Server where jobs are run.

2.2 Using the Orchestrate User Portal

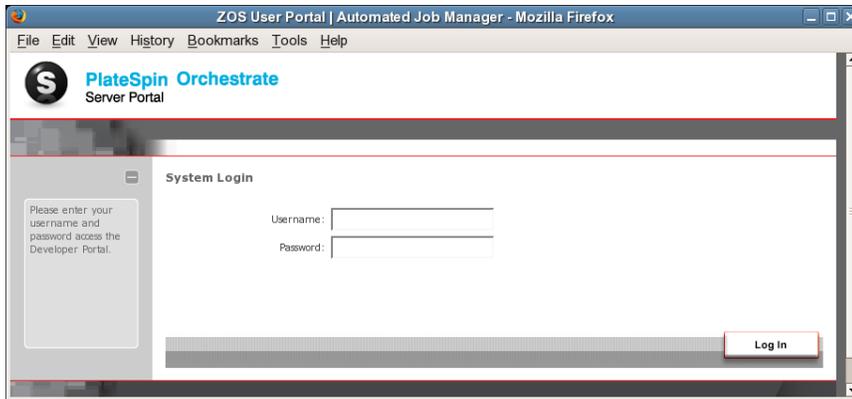
This section describes how to perform some basic tasks with the portal for PlateSpin Orchestrate. The following content is included:

- ♦ [Section 2.2.1, “Connecting to the Orchestrate User Portal,” on page 11](#)
- ♦ [Section 2.2.2, “Displaying Available Jobs,” on page 13](#)
- ♦ [Section 2.2.3, “Running a Job,” on page 14](#)
- ♦ [Section 2.2.4, “Monitoring a Job,” on page 19](#)
- ♦ [Section 2.2.5, “Viewing the Job Log,” on page 20](#)
- ♦ [Section 2.2.6, “Pausing and Resuming a Running Job,” on page 20](#)
- ♦ [Section 2.2.7, “Canceling a Running Job,” on page 21](#)
- ♦ [Section 2.2.8, “Logging Out of the Orchestrate User Portal,” on page 21](#)

2.2.1 Connecting to the Orchestrate User Portal

Use the following steps to run the portal and log in at a workstation:

- 1 Start any Internet browser.
- 2 Enter the URL (server name) of the Orchestrate Server to display a page where the user can log in to the server:

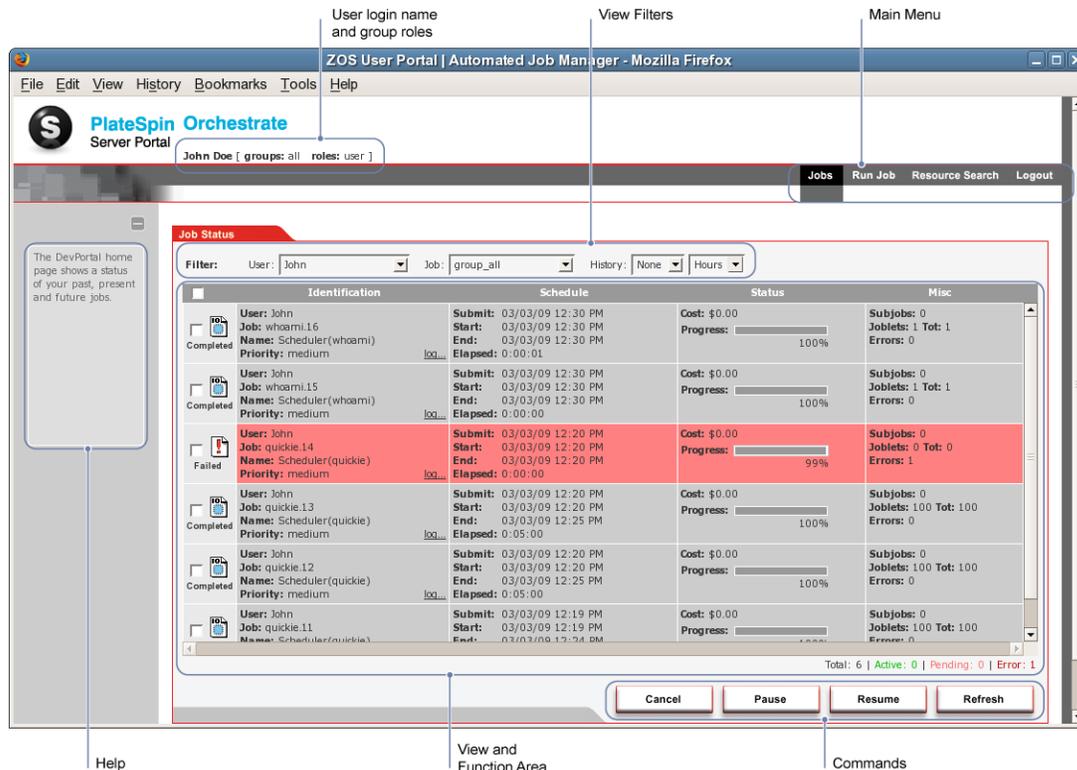


- 3 Specify the username and password created for you by the PlateSpin Orchestrate administrator.
- 4 Click *Log In* to log into the portal.

The Orchestrate User Portal Interface

After you log in to the portal, you see the following interface displayed in the browser:

Figure 2-1 *The User Portal Interface*



NOTE: Logging in with the username of a PlateSpin Orchestrate administrator displays additional menu options for Server Status and User search.

The table below provides details about the various labeled areas of the interface.

Table 2-1 *User Portal Interface Areas*

Area	Description
User Login and Group Roles	Indicates the user groups that the currently logged-in user belongs to and whether the user has an “administrator” role or a “user” role.
Main Menu	Provides operations categorized under the menus: <ul style="list-style-type: none">◆ <i>Jobs</i> include a list of running and completed jobs on the Orchestrate Server and their status.◆ <i>Run Job</i> allows the user to submit a job to the Orchestrate Server, optionally specifying parameter values for the job.◆ <i>Resource Search</i> specifies which resources are to be selected for use by joblets associated with the jobs to be run.◆ <i>Logout</i> returns the user to the login screen.
Filters	Provides a way to filter the view of running and completed jobs.
View and Function	Displays jobs, computing resources, or other requested information. This is the general area of the portal for viewing requested data.
Commands	Provide a set of commands for the functions that can be performed on the data in the View and Function area.
Help	When expanded, displays context sensitive help information when the user mouses over some areas of the interface.

2.2.2 Displaying Available Jobs

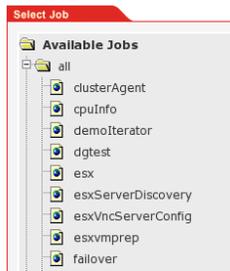
After you log in, you can use the portal to display available jobs that can be run on the PlateSpin Orchestrate system.

To see runnable jobs from the portal command menus, click Run Job in the menu bar to open the Run Job page. The Run Job page allows you to specify what job to run, where the selected job’s joblets execute, and when the job runs.

The Select Job tab is listed in the View and Function area of the portal interface. The Available Jobs explorer is located on the left side of this interface. The Job Parameters area is located on the right side. The Job Parameters area is associated with job selections in the Available Jobs explorer. To view all jobs, expand the explorer under Available Jobs > all.

PlateSpin Orchestrate includes prepackaged jobs for demonstration. Prepackaged jobs are optionally deployed with each new server when it is first started. These jobs include `demoIterator`, `failover`, `quickie`, `subjobs` and so on.

Figure 2-2 The Available Jobs Explorer View Expanded to Show All Jobs



2.2.3 Running a Job

There are four stages to running an PlateSpin Orchestrate job on managed computing resources. This section uses the `quickie` example job to illustrate these stages:

- ♦ “Selecting the Job You Want to Run” on page 14
- ♦ “Selecting Computing Resources to Run the Job” on page 15
- ♦ “Scheduling a Time to Run the Job” on page 17
- ♦ “Running the Job” on page 18

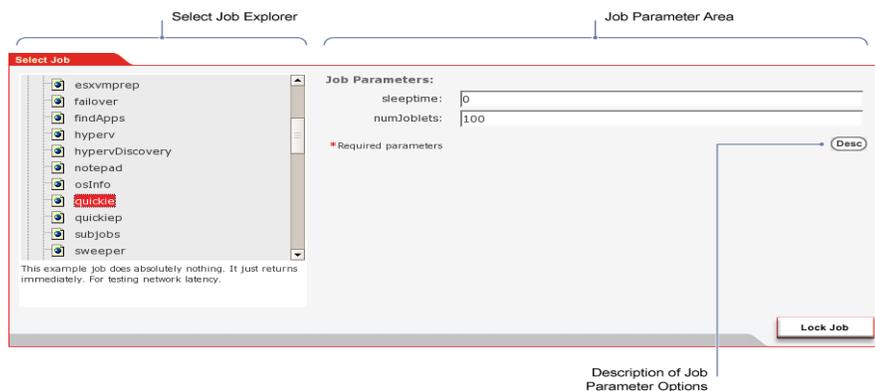
NOTE: Quickie is a demonstration job used primarily to show computing resource utilization. Its purpose is to simply start its executable program on an appropriate computing resource in the grid, wait for a period of time, and then stop to free the computing resource again.

Selecting the Job You Want to Run

When you want to run a job, you need to first select the job and specify its parameters. The portal interface labels this as the *What* stage of running a job.

Use the following steps to select the job and specify job parameters:

- 1 From the *Select Job Explorer*, expand the *examples* folder in the *Available Jobs* directory tree to list the example jobs that are deployed on the PlateSpin Orchestrate Server.



- 2 Select the `quickie` job to display its job parameters in the *Job Parameters* area of the interface. If you need more information about a job parameter, click `Desc` to expand the description of the fields of the *Job Parameters* area.

3 Accept the default values for the parameters.

The `numbjoblets` parameter is an argument for the associated job. The default parameter value for the `quickie` job is 100.

4 Click *Lock Job* to select the job.

The *What* stage on the interface now lists the *quickie* job selection. If you want to change the job selection or its parameters, you can click *Edit* on this stage anytime before you actually run the job.

Selecting Computing Resources to Run the Job

After the job is selected, users need to identify the computing resources where the joblets associated with the job are to be run. The portal interface labels this as the *Where* stage.

Searching for computing resources is done through *Resource Search*. When the Resource Search page is displayed, users can select computing resource search criteria to search for the computing resources they want.

Your grid might have different kinds of computing resources, identified by a set of facts about each computing resource. Examples of resource facts include the resource name, hardware configuration, operating system, and other characteristics. If the grid has uniform resources, the *Where* stage is less useful, but for a grid with multiple resource characteristics, the *Where* stage is very important.

Depending on the resources that you have set up with the PlateSpin Orchestrate Server, you might have multiple computing resource operating system types (operating system fact). For example, you could have both Windows* and Linux operating system resources available. Some of the example jobs that come with the PlateSpin Orchestrate Server have executables for both operating systems, and others do not. The `quickie` job can run on both Windows and Linux.

The following sections contain more information about identifying computing resources:

- ♦ [“Using Expressions to Identify Appropriate Computing Resources” on page 15](#)
- ♦ [“Sample Job: Selecting Computing Resources for Running a Job” on page 16](#)

Using Expressions to Identify Appropriate Computing Resources

A simple expression is used for identifying the computing resources appropriate for running a job. Multiple expressions can be applied. Each expression is entered by using the expression criteria listed in the following table.

Table 2-2 Expressions Used to Filter Computing Resources Eligible to Run Jobs

Expression	Criteria
Fact	Each computing resource has a set of facts. A complete list of facts is available from the <i>Select Fact</i> drop-down list for computing resources.
Operator	Each expression requires one operator.

Expression	Criteria
Value	<p>The value criteria are evaluated against the associated fact according to the selected operator. The portal generates a list of possible value options based on the selected fact and the possible values from the current set of resources.</p> <p>When you select a value, it is entered into the criteria edit field to the right of the value selection field. You can edit the value in this field if necessary.</p>

Multiple criteria can be entered in the same fashion. The selected computing resources match all criteria entered. You can alter your choice of expressions by using the following functions:

- ♦ To remove a criterion, click the minus (-) icon located to the right of the criterion.
- ♦ To start a new search or to start over, click *Reset* located below the criteria list.
- ♦ To add more criteria, click the plus (+) icon.

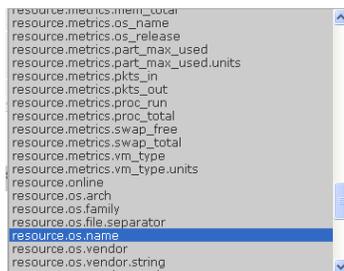
When you have entered all of the search criteria, click *Search*. The matching resources are listed on the lower part of the page under *Search Results*.

Sample Job: Selecting Computing Resources for Running a Job

If no search facts are selected, PlateSpin Orchestrate assumes that all computing resources are to be selected.

The following steps demonstrate how the *quickie* job is configured to run only on Windows XP operating system computing resources:

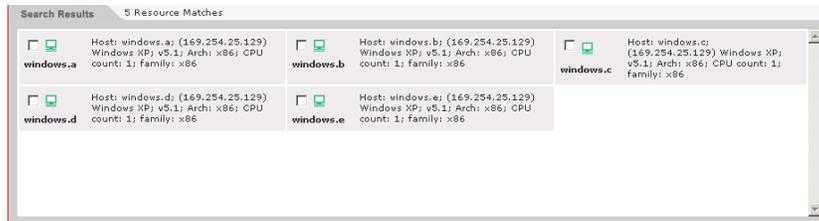
- 1 From the *Select Fact* drop-down list, scroll to and select the operating system fact `resource.os.name`.



- 2 From the *Select Operator* drop-down list, specify the operator to use for the fact. In this case the default, *is equal to*, is used.
- 3 From the *Select Value* drop-down list, specify the fact value, *Windows XP*.

Each of the drop-down items is context sensitive. When a fact is selected, the values update with the valid options for that fact, based on the current set of available computing resources. The values change automatically as computing resources are added or removed in the grid of managed resources.

- 4 Click *Search* to check whether any computing resources match the current search filter criteria. If so, these computing resources are listed at the bottom of the Function area of the portal. See the figure below for an example.



- 5 Select the check box next to the desired resources, then click *Lock Criteria* (located at the bottom right of the Function area of the portal) to select the computing resources.

Notice that the *Where* section (located at the top of the Function area of the portal) now shows the quickie job resource selections. To change the computing resource selection, click *Edit* before running a job.

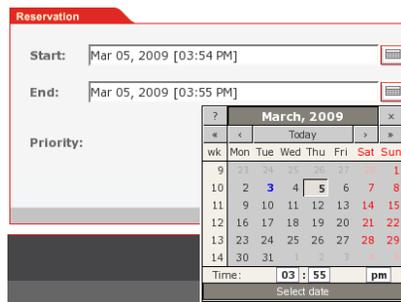
Scheduling a Time to Run the Job

After you have selected a job and the computing resources to be used, you need to select a start time for running the job. The portal labels this as the *When* stage of running a job.

The *Reservation* tab area shows start and end requirements, as well as a priority setting for running the job. In the *When* stage, the user can specify a start and end time and a job priority for the selected job.

Use the following steps to specify when the job is to be run.

- 1 Click the *Reservation* tab, click the calendar icon located adjacent to the *Start* field to open a calendar scheduling utility, then select a date and time to run the job. You can also accept the default in the *Start* field: *As Soon As Possible*.



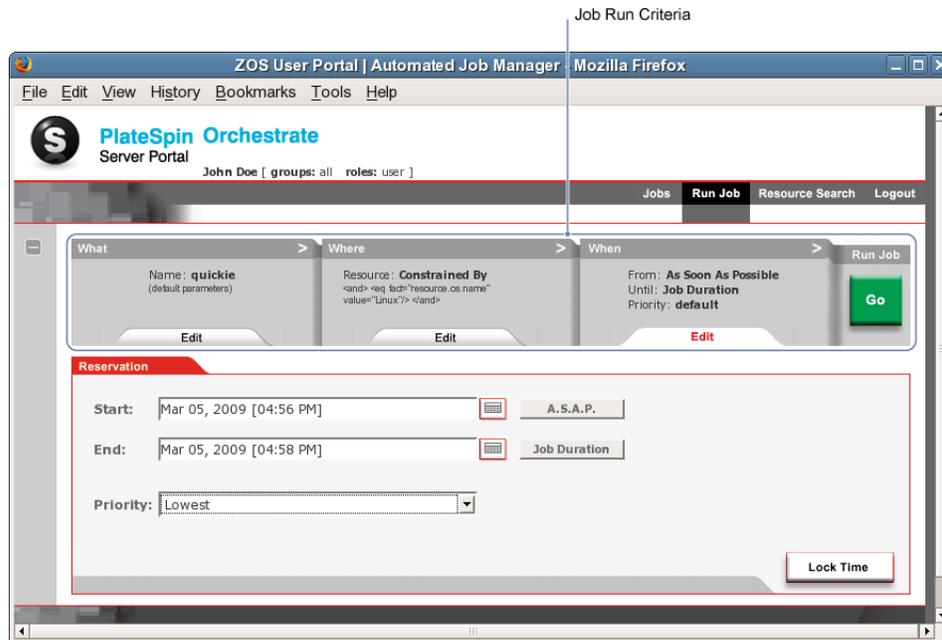
- 2 (Optional) Click the calendar icon located adjacent to the *End* field, then select a date and time to complete the job. You can also accept the default in the *End* field: *Job Duration*.
- 3 (Optional) In the *Priority* field, click the drop-down list to select a priority for the job, or accept the default.

Use the *Priority* parameter to increase or decrease the priority of the job in relation to other jobs that are already running or scheduled to run. This option is used only if computing resources are inadequate to run the job at the scheduled job start time. The options in the drop-down list are context-sensitive and match the available priorities associated with the user profile.

- 4 Click *Lock Time* (located at the bottom right of the Function area) to finalize the time to run the job.

The job is now ready to run. The job criteria are displayed in the upper part of the Function area of the portal.

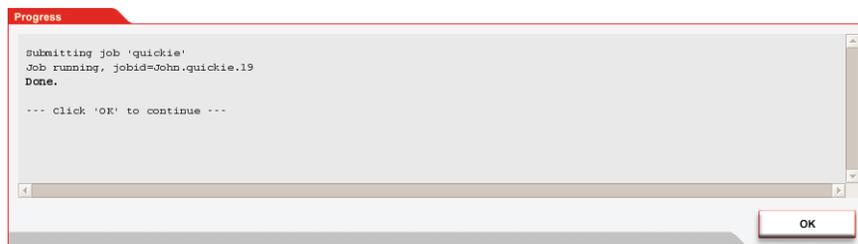
Figure 2-3 The User Portal Displaying the Job Run Criteria Previously Selected



Running the Job

When you have established the job run criteria in the User Portal, click *Go* to schedule the job to run. A progress log is displayed in the *Progress* tab in the Function area of the portal.

Figure 2-4 Progress Message Displayed for the “Quickie” Job



A *Done* message indicates that the job was successfully submitted to the server for PlateSpin Orchestrate.

Click *OK* to close the Progress log and return to the Job Status page to monitor the running job.

NOTE: This same job could also be run using the following *zos* command:

```
zos run quickie.job
```

For more information about using the *zos* command line interface, see “[The zos Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*.

2.2.4 Monitoring a Job

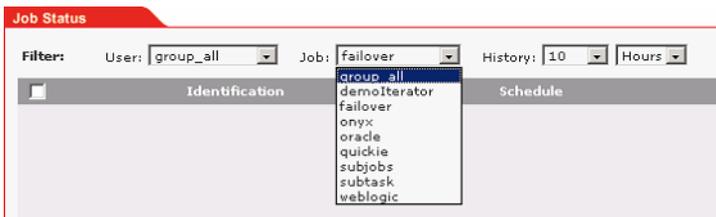
Users can access a list of PlateSpin Orchestrate jobs that are currently running by clicking the *Jobs* command menu.

Figure 2-5 *Jobs Command Menu on the User Portal*



The *Job Status* page lists information about jobs submitted to PlateSpin Orchestrate.

Figure 2-6 *Job Status Page Showing the Job Filter Drop Down List*



A list of jobs is displayed based on the filter criteria selected. The default filter uses all jobs for the currently logged-in user. Filter options include:

- ♦ **Users:** Select a user, or one of the built-in group types: *group_all* for all user groups, or *group_system* for all system jobs.
- ♦ **Job:** Select a specific job, or the built-in group types: *group_all* for all jobs.
- ♦ **History:** Select jobs submitted within the last *nn* hours or *nn* days from the current date and time. The minimum time you can set is 1 hour. The maximum time you can set is 24 days. If you select *none* as the history time, only the most recent jobs are selected.

When you have selected the filter criteria you want, you can update the currently displayed list. Click *Refresh* (located at the bottom right of the interface).

Figure 2-7 *Status Information for a Current Job*

	Identification	Schedule	Status	Misc
	User: paul Job: quickie.645 Name: quickie Priority: medium	Submit: 09/03/05 03:17 AM Start: 09/03/05 03:17 AM End: 09/03/05 03:26 AM Elapsed: 0:09:14	Cost: \$0.05 Progress: <div style="width: 100%;"><div style="width: 100%;"></div></div> 100%	Subjobs: 0 Jobs: 1000 Tot: 1000 Errors: 0

The job status report includes the following columns:

- ♦ The left column indicates the result for the job status report as sampled at refresh time. The column displays the result with both an icon and text indicators. The job result types are *Canceled*, *Scheduled*, *Running*, *Failed*, or *Completed*.
- ♦ **Identification:** This column lists job submission information, including *User*, *Job*, *Name*, and *Priority*.
- ♦ **Schedule:** This column lists information about when the job is scheduled, or its actual start time. It also includes how long the job has been running. If the job is completed, the column also includes the total elapsed run time for job.

- ♦ **Status:** This column lists information about the cost and progress of a job. For jobs that are currently running, the progress indicator bar is partially filled with the percentage of job completion. *Cost* indicates the current cost utilization for a running job, or the total cost for completed and failed jobs. The cost of a job is based on facts configured by administrator for various resources.
- ♦ **Misc:** This column lists other useful information about a job, including its *Subjobs*, its *Joblets* (the number of independent pieces of work running on each computing resource), and any *Errors*.

If the number of subjobs is greater than 0, a *view...* link displays in this column. Click this link to show a Job Tree hierarchy under the *Filter* settings for the job. All of the subjobs in the job hierarchy are shown.

Figure 2-8 Job Tree Hierarchy Displayed on the Job Status Page



Click *Root* in the tree to return to the default Job Status display, which is based on the current filter criteria.

Using the Command Line Interface to Monitor Jobs

Use the following command if you want to use the command line interface to monitor the job:

```
zos status
```

By default, PlateSpin Orchestrate returns the status of the last job to run.

For more detailed information about using the zos command line interface, see “[The zos Command Line Tool](#)” in the *PlateSpin Orchestrate 2.0 Command Line Reference*.

2.2.5 Viewing the Job Log

From the *Jobs* menu in the portal interface, some jobs display a *log...* link on the Identification column of the *Job Status* page.

Figure 2-9 Job Identification with an Available Log Link



Selecting this link displays the logging information for the job. The logging display stays current for the job shown running. When the job is completed, its complete log is displayed.

2.2.6 Pausing and Resuming a Running Job

From the *Jobs* menu in the portal, display the job to pause or resume by including it in the *Job Status* list using the filters. Use the check box on the left side to select the running job to pause or resume. Then click *Pause* or *Resume* on the bottom of the page. A new page displays, indicating the status of the request. Click *OK* to return to the *Job Status* page.

2.2.7 Canceling a Running Job

- 1 From the *Jobs* menu, display the job to cancel by including it in the *Job Status* list using the filters.
- 2 Select the check box for the running job you want to cancel.
- 3 Click *Cancel*.
A new page displays, indicating the status of the cancellation request.
- 4 Click *OK* to return to the *Job Status* page.

2.2.8 Logging Out of the Orchestrate User Portal

To log out of the Orchestrate User Portal, click the *Logout* option under the Main Menu.